

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
3 November 2005 (03.11.2005)

PCT

(10) International Publication Number
WO 2005/103925 A1

(51) International Patent Classification: G06F 15/16

(21) International Application Number:
PCT/AU2005/000579

(22) International Filing Date: 22 April 2005 (22.04.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2004902146 22 April 2004 (22.04.2004) AU

(71) Applicant (for all designated States except US):
WARATEK PTY LIMITED [AU/AU]; Suite 18, 12
Tyron Road, Lindfield, NSW 2070 (AU).

(72) Inventor: and

(75) Inventor/Applicant (for US only): HOLT, John,
Matthew [AU/AU]; Suite 18, 12 Tyron Road, Lindfield,
NSW 2070 (AU).

(74) Agent: FRASER OLD & SOHN; Patent Attorneys, Level
6, 118 Alfred Street, Milsons Point, NSW 2061 (AU).

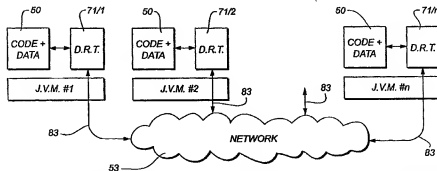
(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,
GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,
KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA,
MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM,
PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY,
TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU,
ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,
FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO,
SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN,
GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— with international search report

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: MULTIPLE COMPUTER ARCHITECTURE WITH SYNCHRONIZATION



(57) Abstract: The present invention discloses a modified computer architecture (50, 71, 72) which enables an applications program (50) to be run simultaneously on a plurality of computers (M1, ... Mn). Shared memory at each computer is updated with amendments and/or overwrites so that all memory read requests are satisfied locally. During initial program loading (75), or similar, instructions which result in the application program (50) acquiring (or releasing) a lock on a particular asset (50A, 50X-50Y) (synchronization) are identified. Additional instructions are inserted (162, 163) to result in a modified synchronization routine with which all computers are updated.

MULTIPLE COMPUTER ARCHITECTURE WITH SYNCHRONIZATION

Field of the Invention

The present invention relates to computers and, in particular, to a modified machine architecture which enables the operation of an application program simultaneously on a plurality of computers interconnected via a communications network.

Background Art

Ever since the advent of computers, and computing, software for computers has been written to be operated upon a single machine. As indicated in Fig. 1, that single prior art machine 1 is made up from a central processing unit, or CPU, 2 which is connected to a memory 3 via a bus 4. Also connected to the bus 4 are various other functional units of the single machine 1 such as a screen 5, keyboard 6 and mouse 7.

A fundamental limit to the performance of the machine 1 is that the data to be manipulated by the CPU 2, and the results of those manipulations, must be moved by the bus 4. The bus 4 suffers from a number of problems including so called bus "queues" formed by units wishing to gain an access to the bus, contention problems, and the like. These problems can, to some extent, be alleviated by various stratagems including cache memory, however, such stratagems invariably increase the administrative overhead of the machine 1.

Naturally, over the years various attempts have been made to increase machine performance. One approach is to use symmetric multiple processors. This prior art approach has been used in so called "super" computers and is schematically indicated in Fig. 2. Here a plurality of CPU's 12 are connected to global memory 13. Again, a bottleneck arises in the communications between the CPU's 12 and the memory 13. This process has been termed "Single System Image". There is only one application and one whole copy of the memory for the application which is distributed over the global memory. The single application can read from and write to, (ie share) any memory location completely transparently.

Where there are a number of such machines interconnected via a network, this is achieved by taking the single application written for a single machine and

partitioning the required memory resources into parts. These parts are then distributed across a number of computers to form the global memory 13 accessible by all CPU's 12. This procedure relies on masking, or hiding, the memory partition from the single running application program. The performance degrades when one CPU on one
5 machine must access (via a network) a memory location physically located in a different machine.

Although super computers have been technically successful in achieving high computational rates, they are not commercially successful in that their inherent complexity makes them extremely expensive not only to manufacture but to
10 administer. In particular, the single system image concept has never been able to scale over "commodity" (or mass produced) computers and networks. In particular, the Single System Image concept has only found practical application on very fast (and hence very expensive) computers interconnected by very fast (and similarly expensive) networks.

A further possibility of increased computer power through the use of a plural number of machines arises from the prior art concept of distributed computing which is schematically illustrated in Fig. 3. In this known arrangement, a single application program (Ap) is partitioned by its author (or another programmer who has become
15 familiar with the application program) into various discrete tasks so as to run upon, say, three machines in which case n in Fig. 3 is the integer 3. The intention here is that each of the machines M1...M3 runs a different third of the entire application and the intention is that the loads applied to the various machines be approximately equal. The machines communicate via a network 14 which can be provided in various forms
20 such as a communications link, the internet, intranets, local area networks, and the like. Typically the speed of operation of such networks 14 is an order of magnitude slower than the speed of operation of the bus 4 in each of the individual machines M1, M2, etc.

Distributed computing suffers from a number of disadvantages. Firstly, it is a difficult job to partition the application and this must be done manually. Secondly,
30 communicating data, partial results, results and the like over the network 14 is an administrative overhead. Thirdly, the need for partitioning makes it extremely difficult to scale upwardly by utilising more machines since the application having

been partitioned into, say three, does not run well upon four machines. Fourthly, in the event that one of the machines should become disabled, the overall performance of the entire system is substantially degraded.

A further prior art arrangement is known as network computing via "clusters" as is schematically illustrated in Fig. 4. In this approach, the entire application is loaded onto each of the machines M1, M2 ...Mn. Each machine communicates with a common database but does not communicate directly with the other machines. Although each machine runs the same application, each machine is doing a different "job" and uses only its own memory. This is somewhat analogous to a number of windows each of which sell train tickets to the public. This approach does operate, is scalable and mainly suffers from the disadvantage that it is difficult to administer the network.

In computer languages such as JAVA and MICROSOFT.NET there are two major types of constructs with which programmers deal. In the JAVA language these are known as objects and classes. In any computer environment it is necessary to acquire and release a lock to enable the use of such assets, resources or structures to avoid different parts of the application program attempting to use the same resource at the one time. In the JAVA environment this is known as synchronization. This is achieved in JAVA by the "monitor enter" and "monitor exit" instructions or routines. Other languages use different terms but utilize a similar concept.

The present invention discloses a computing environment in which an application program operates simultaneously on a plurality of computers. In such an environment it is necessary to ensure that the "monitor enter" and "monitor exit" (more generally synchronization routines) operate in a consistent fashion across all the machines. It is this goal of consistent locking of resources that is the genesis of the present invention.

In accordance with a first aspect of the present invention there is disclosed a method multiple computer system having at least one application program running simultaneously on a plurality of computers interconnected by a communications network, wherein a like plurality of substantially identical objects are created, each in

the corresponding computer and each having a substantially identical name, and said system including a lock means applicable to all said computers wherein any computer wishing to utilize a named object therein acquires an authorizing lock from said lock means which permits said utilization and which prevents all the other computers from
5 utilizing their corresponding named object until said authorizing lock is relinquished.

In accordance with a second aspect of the present invention there is disclosed a plurality of computers interconnected via a communications link and operating at least one application program simultaneously wherein each said computer in operating said
10 at least one application program utilizes an object only in local memory physically located in each said computer, the contents of the local memory utilized by each said computer is fundamentally similar but not, at each instant, identical, and every one of said computers has an acquire lock routine and a release lock routine which permit utilization of the local object only by one computer if each of the remainder of said
15 plurality of computers is locked out of utilization of their corresponding object.

In accordance with a third aspect of the present invention there is disclosed a method of running at least one application program on a plurality of computers simultaneously, said computers being interconnected by means of a communications
20 network, said method comprising the steps of:

- (i) creating a like plurality of substantially identical objects each in the corresponding computer and each having a substantially identical name, and
- (ii) requiring any of said computers wishing to utilize a named object therein to acquire an authorizing lock which permits said utilization and which prevents all the
25 other computers from utilizing their corresponding named object until said authorizing lock is relinquished.

In accordance with a fourth aspect of the present invention there is disclosed a method of ensuring consistent synchronization of an application program to be run simultaneously on a plurality of computers interconnected via a communications
30 network, said method comprising the steps of:

- (i) scrutinizing said application program at, or prior to, or after loading to detect each program step defining a synchronization routine, and
- (ii) modifying said synchronization routine to ensure utilization by only one

computer of an object and preventing all the remaining computers from simultaneously utilizing their corresponding objects.

In accordance with a fifth aspect of the present invention there is disclosed a multiple thread processing computer operation in which individual threads of a single application program are simultaneously being processed each on a corresponding one of a plurality of computers interconnected via a communications link, and in which objects in local memory physically associated with the computer processing each thread have corresponding objects in the local memory of each other said computer, the improvement comprising permitting only one of said computers to utilize an object and preventing all the remaining computers from simultaneously utilizing their corresponding object.

In accordance with a sixth aspect of the present invention there is disclosed a computer program product comprising a set of program instructions stored in a storage medium and operable to permit a plurality of computers to carry out the abovementioned methods.

Brief Description of the Drawings

Embodiments of the present invention will now be described with reference to the drawings in which:

Fig. 1 is a schematic view of the internal architecture of a conventional computer,

Fig. 2 is a schematic illustration showing the internal architecture of known symmetric multiple processors,

Fig. 3 is a schematic representation of prior art distributed computing,

Fig. 4 is a schematic representation of a prior art network computing using clusters,

Fig. 5 is a schematic block diagram of a plurality of machines operating the same application program in accordance with a first embodiment of the present invention,

Fig. 6 is a schematic illustration of a prior art computer arranged to operate JAVA code and thereby constitute a JAVA virtual machine,

Fig. 7 is a drawing similar to Fig. 6 but illustrating the initial loading of code in accordance with the preferred embodiment,

Fig. 8 is a drawing similar to Fig. 5 but illustrating the interconnection of a plurality of computers each operating JAVA code in the manner illustrated in Fig. 7,

Fig. 9 is a flow chart of the procedure followed during loading of the same application on each machine in the network,

5 Fig. 10 is a flow chart showing a modified procedure similar to that of Fig. 9,

Fig. 11 is a schematic representation of multiple thread processing carried out on the machines of Fig. 8 utilizing a first embodiment of memory updating,

Fig. 12 is a schematic representation similar to Fig. 11 but illustrating an alternative embodiment,

10 Fig. 13 illustrates multi-thread memory updating for the computers of Fig. 8,

Fig. 14 is a schematic illustration of a prior art computer arranged to operate in JAVA code and thereby constitute a JAVA virtual machine,

Fig. 15 is a schematic representation of n machines running the application program and serviced by an additional server machine X,

15 Fig. 16 is a flow chart of illustrating the modification of the monitor enter and exit routines,

Fig. 17 is a flow chart illustrating the process followed by processing machine in requesting the acquisition of a lock,

Fig. 18 is a flow chart illustrating the requesting of the release of a lock,

20 Fig. 19 is a flow chart of the response of the server machine X to the request of Fig. 17,

Fig. 20 is a flow chart illustrating the response of the server machine X to the request of Fig. 18,

25 Fig. 21 is a schematic representation of two laptop computers interconnected to simultaneously run a plurality of applications, with both applications running on a single computer,

Fig. 22 is a view similar to Fig. 21 but showing the Fig. 21 apparatus with one application operating on each computer, and

30 Fig. 23 is a view similar to Figs. 21 and 22 but showing the Fig. 21 apparatus with both applications operating simultaneously on both computers.

The specification includes Annexures A and D which provide actual program fragments which implement various aspects of the described embodiments. Annexure A relates to fields and Annexure D to synchronization.

Detailed Description

In connection with Fig. 5, in accordance with a preferred embodiment of the present invention a single application program 50 can be operated simultaneously on a number of machines M1, M2...Mn communicating via network 53. As it will become
5 apparent hereafter, each of the machines M1, M2...Mn operates with the same application program 50 on each machine M1, M2...Mn and thus all of the machines M1, M2...Mn have the same application code and data 50. Similarly, each of the machines M1, M2...Mn operates with the same (or substantially the same) modifier 51 on each machine M1, M2...Mn and thus all of the machines M1, M2...Mn have
10 the same (or substantially the same) modifier 51 with the modifier of machine M2 being designated 51/2. In addition, during the loading of, or preceding the execution of, the application 50 on each machine M1, M2...Mn, each application 50 has been modified by the corresponding modifier 51 according to the same rules (or substantially the same rules since minor optimising changes are permitted within each
15 modifier 51/1 ...51/n).

As a consequence of the above described arrangement, if each of the machines M1, M2...Mn has, say, a shared memory capability of 10MB, then the total shared memory available to each application 50 is not, as one might expect, 10n MB but rather only 10MB. However, how this results in improved operation will become
20 apparent hereafter. Naturally, each machine M1, M2...Mn has an unshared memory capability. The unshared memory capability of the machines M1, M2...Mn are normally approximately equal but need not be.

It is known from the prior art to operate a machine (produced by one of various manufacturers and having an operating system operating in one of various
25 different languages) in a particular language of the application, by creating a virtual machine as schematically illustrated in Fig. 6. The prior art arrangement of Fig. 6 takes the form of the application 50 written in the Java language and executing within a Java Virtual Machine 61. Thus, where the intended language of the application is the language JAVA, a JAVA virtual machine is created which is able to operate code
30 in JAVA irrespective of the machine manufacturer and internal details of the machine. For further details see "The JAVA Virtual Machine Specification" 2nd Edition by T. Lindholm & F. Yellin of Sun Microsystems Inc. of the USA.

This well known prior art arrangement of Fig. 6 is modified in accordance with the preferred embodiment of the present invention by the provision of an additional facility which is conveniently termed "distributed run time" or DRT 71 as seen in Fig. 7. In Fig. 7, the application 50 is loaded onto the Java Virtual Machine
5 72 via the distributed runtime system 71 through the loading procedure indicated by arrow 75. A distributed run time system is available from the Open Software Foundation under the name of Distributed Computing Environment (DCE). In particular, the distributed runtime 71 comes into operation during the loading procedure indicated by arrow 75 of the JAVA application 50 so as to initially create
10 the JAVA virtual machine 72. The sequence of operations during loading will be described hereafter in relation to Fig. 9.

Fig. 8 shows in modified form the arrangement of Fig. 5 utilising JAVA virtual machines, each as illustrated in Fig. 7. It will be apparent that again the same application 50 is loaded onto each machine M1, M2...Mn. However, the
15 communications between each machine M1, M2...Mn, and indicated by arrows 83, although physically routed through the machine hardware, are controlled by the individual DRT's 71/1...71/n within each machine. Thus, in practice this may be conceptualised as the DRT's 71/1...71/n communicating with each other via the network 73 rather than the machines M1, M2...Mn themselves.

Turning now to Figs. 7 and 9, during the loading procedure 75, the program 50 being loaded to create each JAVA virtual machine 72 is modified. This modification commences at 90 in Fig. 9 and involves the initial step 91 of detecting all memory locations (termed fields in JAVA - but equivalent terms are used in other languages) in the application 50 being loaded. Such memory locations need to be identified for
25 subsequent processing at steps 92 and 93. The DRT 71 during the loading procedure 75 creates a list of all the memory locations thus identified, the JAVA fields being listed by object and class. Both volatile and synchronous fields are listed.

The next phase (designated 92 in Fig. 9) of the modification procedure is to search through the executable application code in order to locate every processing
30 activity that manipulates or changes field values corresponding to the list generated at step 91 and thus writes to fields so the value at the corresponding memory location is changed. When such an operation (typically putstatic or putfield in the JAVA

language) is detected which changes the field value, then an "updating propagation routine" is inserted by step 93 at this place in the program to ensure that all other machines are notified that the value of the field has changed. Thereafter, the loading procedure continues in a normal way as indicated by step 94 in Fig. 9.

- 5 An alternative form of initial modification during loading is illustrated in Fig. 10. Here the start and listing steps 90 and 91 and the searching step 92 are the same as in Fig. 9. However, rather than insert the "updating propagation routine" as in step 93 in which the processing thread carries out the updating, instead an "alert routine" is inserted at step 103. The "alert routine" instructs a thread or threads not used in processing and allocated to the DRT, to carry out the necessary propagation.
- 10 This step 103 is a quicker alternative which results in lower overhead.

- Once this initial modification during the loading procedure has taken place, then either one of the multiple thread processing operations illustrated in Figs. 11 and 12 takes place. As seen in Fig. 11, multiple thread processing 110 on the machines
- 15 consisting of threads 111/1...111/4 is occurring and the processing of the second thread 111/2 (in this example) results in that thread 111/2 becoming aware at step 113 of a change of field value. At this stage the normal processing of that thread 111/2 is halted at step 114, and the same thread 111/2 notifies all other machines M2...Mn via the network 53 of the identity of the changed field and the changed value which
- 20 occurred at step 113. At the end of that communication procedure, the thread 111/2 then resumes the processing at step 115 until the next instance where there is a change of field value.

- In the alternative arrangement illustrated in Fig. 12, once a thread 121/2 has become aware of a change of field value at step 113, it instructs DRT processing 120
- 25 (as indicated by step 125 and arrow 127) that another thread(s) 121/1 allocated to the DRT processing 120 is to propagate in accordance with step 128 via the network 53 to all other machines M2...Mn the identity of the changed field and the changed value detected at step 113. This is an operation which can be carried out quickly and thus the processing of the initial thread 111/2 is only interrupted momentarily as indicated
- 30 in step 125 before the thread 111/2 resumes processing in step 115. The other thread 121/1 which has been notified of the change (as indicated by arrow 127) then

communicates that change as indicated in step 128 via the network 53 to each of the other machines M2...Mn.

5 This second arrangement of Fig. 12 makes better utilisation of the processing power of the various threads 111/1...111/3 and 121/1 (which are not, in general, subject to equal demands) and gives better scaling with increasing size of "n", (n being an integer greater than or equal to 2 which represents the total number of machines which are connected to the network 53 and which run the application program 50 simultaneously). Irrespective of which arrangement is used, the changed field and identities and values detected at step 113 are propagated to all the other
10 machines M2...Mn on the network.

This is illustrated in Fig. 13 where the DRT 71/1 and its thread 121/1 of Fig. 12 (represented by step 128 in Fig. 13) sends via the network 53 the identity and changed value of the listed memory location generated at step 113 of Fig. 12 by processing in machine M1, to each of the other machines M2...Mn.

15 Each of the other machines M2...Mn carries out the action indicated by steps 135 and 136 in Fig. 13 for machine Mn by receiving the identity and value pair from the network 53 and writing the new value into the local corresponding memory location.

20 In the prior art arrangement in Fig. 3 utilising distributed software, memory accesses from one machine's software to memory physically located on another machine are permitted by the network interconnecting the machines. However, such memory accesses can result in delays in processing of the order of $10^6 - 10^7$ cycles of the central processing unit of the machine. This in large part accounts for the diminished performance of the multiple interconnected machines.

25 However, in the present arrangement as described above in connection with Fig. 8, it will be appreciated that all reading of data is satisfied locally because the current value of all fields is stored on the machine carrying out the processing which generates the demand to read memory. Such local processing can be satisfied within $10^2 - 10^3$ cycles of the central processing unit. Thus, in practice, there is
30 substantially no waiting for memory accesses which involves reads.

However, most application software reads memory frequently but writes to memory relatively infrequently. As a consequence, the rate at which memory is being written or re-written is relatively slow compared to the rate at which memory is being read. Because of this slow demand for writing or re-writing of memory, the fields can
5 be continually updated at a relatively low speed via the inexpensive commodity network 53, yet this low speed is sufficient to meet the application program's demand for writing to memory. The result is that the performance of the Fig. 8 arrangement is vastly superior to that of Fig. 3.

10 In a further modification in relation to the above, the identities and values of changed fields can be grouped into batches so as to further reduce the demands on the communication speed of the network 53 interconnecting the various machines.

It will also be apparent to those skilled in the art that in a table created by each DRT 71 when initially recording the fields, for each field there is a name or identity which is common throughout the network and which the network recognises.
15 However, in the individual machines the memory location corresponding to a given named field will vary over time since each machine will progressively store changed field values at different locations according to its own internal processes. Thus the table in each of the DRTs will have, in general, different memory locations but each global "field name" will have the same "field value" stored in the different memory
20 locations.

It will also be apparent to those skilled in the art that the abovementioned modification of the application program during loading can be accomplished in up to five ways by:

- (i) re-compilation at loading,
- 25 (ii) by a pre-compilation procedure prior to loading,
- (iii) compilation prior to loading,
- (iv) a "just-in-time" compilation, or
- (v) re-compilation after loading (but, or for example, before execution of the relevant or corresponding application code in a distributed environment).

30 Traditionally the term "compilation" implies a change in code or language, eg from source to object code or one language to another. Clearly the use of the term

"compilation" (and its grammatical equivalents) in the present specification is not so restricted and can also include or embrace modifications within the same code or language.

5 In the first embodiment, a particular machine, say machine M2, loads the application code on itself, modifies it, and then loads each of the other machines M1, M3 ... Mn (either sequentially or simultaneously) with the modified code. In this arrangement, which may be termed "master/slave", each of machines M1, M3, ... Mn loads what it is given by machine M2.

10 In a still further embodiment, each machine receives the application code, but modifies it and loads the modified code on that machine. This enables the modification carried out by each machine to be slightly different being optimized based upon its architecture and operating system, yet still coherent with all other similar modifications.

15

In a further arrangement, a particular machine, say M1, loads the unmodified code and all other machines M2, M3 ... Mn do a modification to delete the original application code and load the modified version.

20 In all instances, the supply can be branched (ie M2 supplies each of M1, M3, M4, etc directly) or cascaded or sequential (ie M2 applies M1 which then supplies M3 which then supplies M4, and so on).

25 In a still further arrangement, the machines M1 to Mn, can send all load requests to an additional machine (not illustrated) which is not running the application program, which performs the modification via any of the aforementioned methods, and returns the modified routine to each of the machines M1 to Mn which then load the modified routine locally. In this arrangement, machines M1 to Mn forward all load requests to this additional machine which returns a modified routine to each
30 machine. The modifications performed by this additional machine can include any of the modifications covered under the scope of the present invention.

Persons skilled in the computing arts will be aware of at least four techniques used in creating modifications in computer code. The first is to make the modification in the original (source) language. The second is to convert the original code (in say JAVA) into an intermediate representation (or intermediate language). Once this
5 conversion takes place the modification is made and then the conversion is reversed. This gives the desired result of modified JAVA code.

The third possibility is to convert to machine code (either directly or via the abovementioned intermediate language). Then the machine code is modified before
10 being loaded and executed. The fourth possibility is to convert the original code to an intermediate representation, which is then modified and subsequently converted into machine code.

The present invention encompasses all four modification routes and also a
15 combination of two, three or even all four, of such routes.

Turning now to Fig. 14, there is illustrated a schematic representation of a single prior art computer operated as a JAVA virtual machine. In this way, a machine (produced by any one of various manufacturers and having an operating system operating in any one of various different languages) can operate in the particular
20 language of the application program 50, in this instance the JAVA language. That is, a JAVA virtual machine 72 is able to operate code 50 in the JAVA language, and utilize the JAVA architecture irrespective of the machine manufacturer and the internal details of the machine.

Furthermore, the single machine of Fig. 14 is able to easily perform
25 synchronization of specific objects 50X-50Z when specified by the programmer's use of a synchronization routine. As each object exists only locally, the single JAVA virtual machine 72 of Fig. 14 is able to ensure that an object is properly synchronized as specified by the programmer and thus only utilized by one part of the executable
30 code at any single point in time. If another part of the executable code wishes to use the same object then the possible contention is resolved by the JAVA virtual machine 72 such that other executing parts of the application program have to wait until the first part has finished.

The same procedure applies *mutatis mutandis* for classes 50A. In particular, the computer programmer when writing a program using the JAVA language and architecture, need only to use a synchronization routine(s) in order to provide for this avoidance of contention. Thus a single JAVA virtual machine can keep track of utilization of the classes and objects and avoid any corresponding problems as necessary in an unobtrusive fashion. The process whereby only one object or class is exclusively used is termed "synchronization". In the JAVA language the instructions "monitorenter" and "monitorexit" signify the beginning and ending of a synchronization routine which results in the acquiring of and releasing of a "lock" which prevents an asset being the subject of contention.

However, in the arrangement illustrated in Fig. 8, (and also in Figs. 20-22), a plurality of individual computers or machines M1, M2 Mn are provided each of which are interconnected via a communications network 53 and each of which is provided with a modifier 51 (as in Fig. 5 and realized by the DRT 71 in Fig. 8) and loaded with a common application program 50. Essentially the modifier 51 or DRT 71 ensures that when part of the application program 50 running on one of the machines exclusively utilizes (eg, by means of synchronization) a particular local asset, such as an objects 50X-50Z or class 50A, no other machine M2...Mn utilizes the corresponding asset in its local memory.

In particular, whilst one particular machine (say, M3) is exclusively using an object or class, another machine (say M5) may also be instructed by the code it is executing to exclusively use that object or class at that time. Thus if the object or class were to be exclusively used by both machines, then the behaviour of the object and application as a whole is undefined – that is, in the absence of proper exclusive use of an object when explicitly specified by the programmer, permanent inconsistency between machine M5 and machine M3 is likely to result. Thus the goal of substantially identical memory contents for each of the machines M1, M2...Mn, as required for simultaneous operation of the same application program, would not be achieved.

In order to ensure consistent synchronization the application program is scrutinized in order to detect program steps which define a synchronization routine. This scrutiny can take place either prior to loading, or during the loading procedure,

- or even after the loading procedure (but before execution of the relevant corresponding portion of the application code). It may be likened to a compilation procedure with the understanding that the term compilation normally involves a change in code or language, eg from source to object code or one language to another.
- 5 However, in the present instance the term "compilation" (and its grammatical equivalents) is not so restricted and can also include embrace modifications within the same code or language.

Reference is made to the accompanying Annexure D in which:

- 10 Annexure D1 is a typical code fragment from an unmodified synchronization routine, and
Annexure D2 is an equivalent in respect of a modified synchronization routine,

- Annexures D1 and D2 are the before and after excerpt of a synchronization
15 routine respectively. The modified code that is added to the method is highlighted in bold. In the original code sample of Annexure D1, the code increments a shared memory location (counter) within in synchronize statement. The purpose of the synchronize statement is to ensure thread-safety of the increment operation in multi-threaded applications. Thus, without management of synchronization in a distributed
20 environment, each machine would perform synchronization in isolation, thus potentially incrementing the shared counter at the same time, leading to potential race condition(s) and incoherent memory. Clearly this is not what the programmer of the application program expects to happen.

- 25 So, taking advantage of the DRT, the application code is modified as it is loaded into the machine by changing the synchronization routine. The changes made (highlighted in bold) are the initial instructions and ending instructions that the synchronization routine executes. These added instructions act to additionally perform synchronization across all other machines in the distributed environment,
30 thereby preserving the synchronize behaviour of the application program across a plurality of machines.

The acquireLock() method of the DRT takes an argument which represents a unique identifier for this object (See Annexure D2), for example the name of the object, a reference to the object in question, or a unique number representing this object across all nodes, to be used in acquiring a global lock of the specified object.

- 5 This way, the DRT can support the synchronization of multiple objects at the same time without becoming confused as to which of the multiple objects are already synchronized and which are not, by using the unique identifier of each object to consult the correct record in the shared synchronization table.

- 10 The DRT can determine the synchronization state of the object in a number of ways. Preferably, it can ask each machine in turn if their local copy of this object is presently synchronized, and if any machine replies true, then to wait until that object is unsynchronised, otherwise synchronize this object locally. Alternatively, the DRT on the local machine can consult a shared record table (perhaps on a separate machine
15 (eg machine X), or a coherent shared record table on the local machine, or a database) to determine if this object has been marked as synchronized by any other machine, and if so, then wait until the status of the object is changed to "unsynchronised" and then acquire the lock by marking the object as synchronized, otherwise acquire the lock by marking the object as synchronized by this machine.

- 20 If the DRT determines that no other machine currently has a lock for this object (ie, no other machine has synchronized this object), then to acquire the lock for this object on all other machines, for example by means of modifying the corresponding entry in a shared table of synchronization states, or alternatively,
25 sequentially acquiring the lock on all other machines in addition the current machine. Only once this machine has successfully confirmed that no other machine has currently synchronized this object, and this machine has correspondingly synchronized locally, can the execution of the original synchronized code-block begin..

30

On the other hand, if the DRT determines that another machine has already synchronized this object, then this machine is to postpone execution of the original
synchronize code-block until such a time as the DRT can confirm that no other

machine is presently executing a synchronize statement for this object, and that this machine has correspondingly synchronized the object locally. In such a case, the original code block is **NOT** to be executed until this machine can guarantee that no other machine is executing a synchronize statement for this object, as it will potentially corrupt the object across the participating machines due to race-conditions, inconsistency of memory, and so forth resulting from the concurrent execution of synchronized statements. Thus, when the DRT determines that this object is presently "synchronized", the DRT prevents execution of the original code-block by pausing the execution of the "acquireLock()" operation until such a time as a corresponding "releaseLock()" operation is executed by the present owner of the lock.

Thus, on execution of a "releaseLock()" operation, the machine which presently "owns" a lock (ie, is executing a synchronized statement) indicates the close of its synchronized statement, for example by marking this object as "unsynchronised" in the shared table of synchronization states, or alternatively, sequentially releasing locks acquired on all other machines. At this point, any other machine waiting to begin execution of a corresponding synchronized statement can then claim ownership of this object's lock by resuming execution of its postponed (ie delayed) "acquireLock()" operation, for example, marking itself as executing a synchronized statement for this object in the shared table of synchronization states, or alternatively, sequentially acquiring local locks on each of the other machines.

So, taking advantage of the DRT, the application code is modified as it is loaded into the machine by changing the synchronization routine (consisting of a beginning "monitorenter" and an ending "monitorexit" instruction/s). The changes made (highlighted in **bold**) are the initial instructions that the synchronization routine executes. These added instructions check if this lock has already been acquired by another machine. If this lock has not been acquired by another machine, then the DRT of this machine notifies all other machines that this machine has acquired the lock, and thereby stopping the other machines from executing synchronization routines for this lock.

The DRT can record the lock status of the machines in many ways, for example:

1. corresponding to the entry to a synchronization routine, the DRT individually consults each machine to ascertain if this lock is already acquired. If so, the DRT
5 pauses the execution of the synchronization routine until all other machines no longer own a lock on this asset or object. Otherwise, the DRT executes this synchronization routine. Alternatively,
2. corresponding to the entry to a synchronization routine, the DRT consults a shared table of records (for example a shared database, or a copy of a shared table on
10 each of the participating machines) which indicate if any machine currently "owns" this lock. If so, the DRT then pauses execution of the synchronization routine on this machine until all other machines no longer own a lock on this object. Otherwise the DRT records this machine in the shared table (or tables, if there are multiple tables of records, eg, on multiple machines) as the owner of this lock, and then executes the
15 synchronization routine.

Similarly, when a lock is released, that is to say, when the execution of a synchronization routine is to end, the DRT can "un-record" the lock status of machines in many alternative ways, for example:

- 20 1. corresponding to the exit to a synchronization routine, the DRT individually notifies each other machine that it no longer owns the lock. Alternatively,
2. corresponding to the exit to a synchronization routine, the DRT updates the record for this locked asset or object in the shared table(s) of records such that this machine is no longer recorded as owning this lock.

25

Still further, the DRT can queue machines needing to acquire a locked object in multiple alternative ways, for example:

1. corresponding to the entry to a synchronization routine, the DRT notifies the present owner of the locked object that a specific machine would like to acquire the
30 lock upon release by the current owning machine. The specified machine, if there are no other waiting machines, then stores a record of the specified machine's interest in a table, which, following the exit of the synchronization routine of the locked object,

then notifies the waiting machine that it can acquire this locked object, and thus begin executing its synchronization routine,

2. corresponding to the entry to a synchronization routine, the DRT notifies the present owner of the locked object that a specific machine (say machine M6) would like to acquire the lock upon release by that machine (say machine M4). That machine M4, if it finds after consulting its records of waiting machines for this locked object, finds that there are already one or more machines waiting, then either appends machine M6 to the end of the list of machines wanting to acquire this locked object, or alternatively, forwards the request from M6 to the first waiting, or any other machine waiting, machine which then, in turn, records machine M6 in their table of records,
3. corresponding to the entry to a synchronization routine, the DRT records itself in a shared table(s) of records (for example, a table stored in a shared database accessible by all machines, or multiple separate tables which are substantially similar).

Still further, the DRT can notify other machines queued to acquire this lock corresponding to the exit of a synchronization routine by this machine in the following alternative ways, for example:

1. corresponding to the exit of a synchronization routine, the DRT notifies one of the awaiting machines (for example, this first machine in the queue of waiting machines) that the lock is released,
2. corresponding to the exit of a synchronization routine, the DRT notifies one of the awaiting machines (for example, the first machine in the queue of waiting machines) that the lock is released, and additionally, provides a copy of the entire queue of machines (for example, the second machine and subsequent machines awaiting for this lock). This way, the second machine inherits the list of waiting machines from the first machine, and thereby ensures the continuity of the queue of waiting machines as each machine in turn down the list acquires and subsequently releases the lock.

During the abovementioned scrutiny, "monitorenter" and "monitorexit" instructions (or methods) are initially looked for and, when found, a modifying code is inserted so as to give rise to a modified synchronization routine. This modified

routine acquires and releases a lock. There are several different modes whereby this modification and loading can be carried out.

As seen in Fig. 15 a modification to the general arrangement of Fig. 8 is provided in that machines M1, M2...Mn are as before and run the same application program 50 (or programmes) on all machines simultaneously. However, the previous arrangement is modified by the provision of a server machine X which is conveniently able to supply housekeeping functions, for example, and especially the synchronization of structures, assets and resources. Such a server machine X can be a low value commodity computer such as a PC since its computational load is low. As indicated by broken lines in Fig. 15, two server machines X and X+1 can be provided for redundancy purposes to increase the overall reliability of the system. Where two such server machines X and X+1 are provided, they are preferably operated as dual machines in a cluster.

It is not necessary to provide a server machine X as its computational load can be distributed over machines M1, M2...Mn. Alternatively, a database operated by one machine (in a master/slave type operation) can be used for the housekeeping function(s).

Fig. 16 shows a preferred general procedure to be followed. After loading has been commenced, the instructions to be executed are considered in sequence and all synchronization routines are detected as indicated in step 162. In the JAVA language these are the "monitorenter" and "monitorexit" instructions. Other languages use different terms.

Where a synchronization routine is detected, it is modified, typically by inserting further instructions into the routine. Alternatively, the modifying instructions could be inserted prior to the routine. Once the modification has been completed the loading procedure continues. The modifications preferably take the form of an "acquire lock on all other machines" operation and a "release lock on all other machines" modification as indicated at step 163.

Fig. 17 illustrates a particular form of modification. Firstly, the structures, assets or resources (in JAVA termed classes or objects eg 50A, 50X-50Y) to be

synchronized have already been allocated a name or tag which can be used globally by all machines, as indicated by step 172. This preferably happens when the classes or objects are originally initialized. This is most conveniently done via a table maintained by server machine X. This table also includes the synchronization status
5 of the class or object. In the preferred embodiment, this table also includes a queue arrangement which stores the identities of machines which have requested use of this asset.

As indicated in step 173 of Fig. 17, next an "acquire lock" request is sent to machine X, after which, the sending machine awaits for confirmation of lock
10 acquisition as shown in step 174. Thus, if the global name is already locked (ie the corresponding asset is in use by another machine other than the machine proposing to acquire the lock) then this means that the proposed synchronization routine of the object or class should be paused until the object or class is unlocked by the current owner.

15 Alternatively, if the global name is not locked, this means that no other machine is using this class or object, and confirmation of lock acquisition is received straight away. After receipt of confirmation of lock acquisition, execution of the synchronization routine is allowed to continue, as shown in step 175

Fig. 18 shows the procedures followed by the application program executing
20 machine which wishes to relinquish a lock. The initial step is indicated at step 181. The operation of this proposing machine is temporarily interrupted by steps 183, 184 until the reply is received from machine X, corresponding to step 184, and execution then resumes as indicated in step 185. Optionally, and as indicated in step 182, the machine requesting release of a lock is made to lookup the "global name" for this lock
25 preceding a request being made to machine X. This way, multiple locks on multiple machines can be acquired and released without interfering with one another.

Fig. 19 shows the activity carried out by machine X in response to an "acquire lock" enquiry (of Fig. 17). After receiving an "acquire lock" request at step 191, the lock status is determined at steps 192 and 193 and, if no - the named resource is not
30 free, the identity of the enquiring machine is added at step 194 to (or forms) the queue of awaiting acquisition requests. Alternatively, if the answer is yes - the named

resource is free- the corresponding reply is sent at step 197. The waiting enquiring machine is then able to execute the synchronization routine accordingly by carrying out step 175 of Fig. 17. In addition to the yes response, the shared table is updated at step 196 so that the status of the globally named asset is changed to "locked".

5 Fig. 20 shows the activity carried out by machine X in response to a "release lock" request of Fig. 18. After receiving a "release lock" request at step 201, machine X optionally, and preferably, confirms that the machine requesting to release the lock is indeed the current owner of the lock", as indicated in step 202. Next, the queue status is determined at step 203 and, if no-one is waiting to acquire this lock, machine
10 X marks this lock as "unowned" in the shared table, as shown in step 207, and optionally sends a confirmation of release back to the requesting machine, as indicated by step 208. This enables the requesting machine to execute step 185 of Fig. 18.

 Alternatively, if yes – that is, other machines are waiting to acquire this lock -
15 machine X marks this lock as now acquired by the next machine in the queue, as shown in step 204, and then sends a confirmation of lock acquisition to the queued machine at step 205, and consequently removes the new lock owner from the queue of waiting machines, as indicated in step 206.

 Given the fundamental concept of modifying the synchronization routines
20 there are several different ways in which this concept can be implemented.

 In the first embodiment, a particular machine, say machine M2, loads the synchronization routine on itself, modifies it, and then loads each of the other machines M1, M3 ... Mn (either sequentially or simultaneously) with the modified
25 synchronization routine. In this arrangement, which may be termed "master/slave" each of machines M1, M3, ... Mn loads what it is given by machine M2.

 In a variation of this "master/slave" arrangement, machine M2 loads the synchronization routine in unmodified form on machine M2 and then on the other
30 machines deletes the synchronization routine in its entirety and loads the modified code. Thus in this instance the modification is not a by-passing of the synchronization routine but a deletion of it on all machines except one.

In a still further embodiment, each machine receives the synchronization routine, but modifies it and loads the modified routine on that machine. This enables the modification carried out by each machine to be slightly different being optimized
5 based upon its architecture and operating system, yet still coherent with all other similar modifications.

In a further arrangement, a particular machine, say M1, loads the unmodified synchronization routine and all other machines M2, M3 ... Mn do a modification to
10 delete the original synchronization routine and load the modified version.

In all instances, the supply can be branched (ie M2 supplies each of M1, M3, M4, etc directly) or cascaded or sequential (ie M2 applies M1 which then supplies M3 which then supplies M4, and so on).

In a still further arrangement, the machines M1 to Mn, can send all load requests to an additional machine X (of Fig. 15), which performs the modification via any of the afore mentioned methods, and returns the modified routine to each of the machines M1 to Mn which then load the modified routine locally. In this arrangement, machines M1 to Mn forward all load requests to machine X, which
15 returns a modified routine to each machine. The modifications performed by machine X can include any of the modifications covered under the scope of the present invention.
20

Persons skilled in the computing arts will be aware of four techniques used in
25 creating modifications in computer code. The first is to make the modification in the original (source) language. The second is to convert the original code (in say JAVA) into an intermediate representation (or intermediate language). Once this conversion takes place the modification is made and then the conversion is reversed. This gives the desired result of modified JAVA code.

30

The third possibility is to convert to machine code (either directly or via the abovementioned intermediate language). Then the machine code is modified before being loaded and executed. The fourth possibility is to convert the original code to an

intermediate representation, which is then modified and subsequently converted into machine code.

The present invention encompasses all four modification routes and also a
5 combination of two, three or even all four, of such routes.

Turning now to Figs. 21-23, two laptop computers 101 and 102 are illustrated. The computers 101 and 102 are not necessarily identical and indeed, one can be an IBM or IBM-clone and the other can be an APPLE computer. The computers 101 and 102 have two screens 105, 115 two keyboards 106, 116 but a single mouse 107. The
10 two machines 101, 102 are interconnected by a means of a single coaxial cable or twisted pair cable 314.

Two simple application programs are downloaded onto each of the machines 101, 102, the programs being modified as they are being loaded as described above. In this embodiment the first application is a simple calculator program and results in
15 the image of a calculator 108 being displayed on the screen 105. The second program is a graphics program which displays four coloured blocks 109 which are of different colours and which move about at random within a rectangular box 310. Again, after loading, the box 310 is displayed on the screen 105. Each application operates independently so that the blocks 109 are in random motion on the screen 105 whilst
20 numerals within the calculator 108 can be selected (with the mouse 107) together with a mathematical operator (such as addition or multiplication) so that the calculator 108 displays the result.

The mouse 107 can be used to "grab" the box 310 and move same to the right across the screen 105 and onto the screen 115 so as to arrive at the situation illustrated
25 in Fig. 22. In this arrangement, the calculator application is being conducted on machine 101 whilst the graphics application resulting in display of box 310 is being conducted on machine 102.

However, as illustrated in Fig. 23, it is possible by means of the mouse 107 to drag the calculator 108 to the right as seen in Fig. 22 so as to have a part of the
30 calculator 108 displayed by each of the screens 105, 115. Similarly, the box 310 can be dragged by means of the mouse 107 to the left as seen in Fig. 22 so that the box

310 is partially displayed by each of the screens 105, 115 as indicated Fig. 23. In this configuration, part of the calculator operation is being performed on machine 101 and part on machine 102 whilst part of the graphics application is being carried out the machine 101 and the remainder is carried out on machine 102.

- 5 The foregoing describes only some embodiments of the present invention and modifications, obvious to those skilled in the art, can be made thereto without departing from the scope of the present invention. For example, reference to JAVA includes both the JAVA language and also JAVA platform and architecture.

- 10 Those skilled in the programming arts will be aware that when additional code or instructions is/are inserted into an existing code or instruction set to modify same, the existing code or instruction set may well require further modification (eg by re-numbering of sequential instructions) so that offsets, branching, attributes, mark up and the like are catered for.

- 15 Similarly, in the JAVA language memory locations include, for example, both fields and array types. The above description deals with fields and the changes required for array types are essentially the same mutatis mutandis. Also the present invention is equally applicable to similar programming languages (including procedural, declarative and object orientated) to JAVA including Microsoft.NET platform and architecture (Visual Basic, Visual C/C++, and C#) FORTRAN, C/C++,
20 COBOL, BASIC etc.

- The abovementioned embodiment in which the code of the JAVA synchronization routine is modified, is based upon the assumption that either the run time system (say, JAVA HOTSPOT VIRTUAL MACHINE written in C and JAVA) or the operating system (LINUX written in C and Assembler, for example) of each
25 machine M1...Mn will normally acquire the lock on the local machine (say M2) but not on any other machines (M1, M3 ...Mn). It is possible to leave the JAVA synchronization routine unamended and instead amend the LINUX or HOTSPOT routine which acquires the lock locally, so that it correspondingly acquires the lock on all other machines as well. In order to embrace such an arrangement the term
30 "synchronization routine" is to be understood to include within its scope both the JAVA synchronization routine and the "combination" of the JAVA synchronization

routine and the LINUX or HOTSPOT code fragments which perform lock acquisition and release.

- The terms object and class used herein are derived from the JAVA environment and are intended to embrace similar terms derived from different environments such as dynamically linked libraries (DLL), or object code packages, or function unit or memory locations.

The term "comprising" (and its grammatical variations) as used herein is used in the inclusive sense of "having" or "including" and not in the exclusive sense of "consisting only of".

10 Copyright Notice

- This patent specification contains material which is subject to copyright protection. The copyright owner (which is the applicant) has no objection to the reproduction of this patent specification or related materials from publicly available associated Patent Office files for the purposes of review, but otherwise reserves all copyright whatsoever. In particular, the various instructions are not to be entered into a computer without the specific written approval of the copyright owner.

Annexure A

The following are program listings in the JAVA language:

A1. This first excerpt is part of the modification code. It searches through the code array, and when it finds a putstatic instruction (opcode 178), it implements the modifications.

```
// START
byte[] code = Code_attribute.code; // Bytecode of a given method in a
// given classfile.

int code_length = Code_attribute.code_length;

int DRT = 99; // Location of the CONSTANT_Methodref_info for the
// DRT.alert() method.

for (int i=0; i<code_length; i++){

    if ((code[i] & 0xff) == 179){ // Putstatic instruction.

        System.arraycopy(code, i+3, code, i+6, code_length-(i+3));

        code[i+3] = (byte) 184; // Invokestatic instruction for the
        // DRT.alert() method.

        code[i+4] = (byte) ((DRT >> 8) & 0xff);

        code[i+5] = (byte) (DRT & 0xff);

    }

}

// END
```

A2. This second excerpt is part of the DRT.alert() method. This is the body of the DRT.alert() method when it is called.

```
// START
public static void alert(){

    synchronized (ALERT_LOCK){

        ALERT_LOCK.notify(); // Alerts a waiting DRT thread in the background.

    }

}

// END
```

A3. This third excerpt is part of the DRT Sending. This code fragment shows the DRT in a separate thread, after being notified, sending the value across the network.

```
// START
MulticastSocket ms = DRT.getMulticastSocket(); // The multicast socket
// used by the DRT for
// communication.

byte nameTag = 33; // This is the "name tag" on the network for this
// field.

Field field = modifiedClass.getDeclaredField("myField1"); // Stores
```

```

// the field
// from the
// modified
// class.

// In this example, the field is a byte field.
while (DRT.isRunning()){

    synchronized (ALERT_LOCK){

        ALERT_LOCK.wait(); // The DRT thread is waiting for the alert
                           // method to be called.

        byte[] b = new byte[]{nameTag, field.getBytes(null)}; // Stores
                                                                // the
                                                                // nameTag
                                                                // and the
                                                                // value
                                                                // of the
                                                                // field from
                                                                // the modified
                                                                // class in a
                                                                // buffer.

        DatagramPacket dp = new DatagramPacket(b, 0, b.length);

        ms.send(dp); // Send the buffer out across the network.

    }

}
// END

```

A4. The fourth excerpt is part of the DRT receiving. This is a fragment of code to receive a DRT sent alert over the network.

```

// START
MulticastSocket ms = DRT.getMulticastSocket(); // The multicast socket
                                                // used by the DRT for
                                                // communication.

DatagramPacket dp = new DatagramPacket(new byte[2], 0, 2);

byte nameTag = 33; // This is the "name tag" on the network for this
                  // field.

Field field = modifiedClass.getDeclaredField("myField1"); // Stores the
                                                            // field from
                                                            // the modified
                                                            // class.

// In this example, the field is a byte field.
while (DRT.isRunning){

    ms.receive(dp); // Receive the previously sent buffer from the network.

    byte[] b = dp.getData();

    if (b[0] == nameTag){ // Check the nametags match.
        field.setByte(null, b[1]); // Write the value from the network packet
                                   // into the field location in memory.
    }

}

}
// END

```

A5. The fifth excerpt is an example application before modification has occurred.

```
Method void setValues(int, int)
0 iload_1
1 putstatic #3 <Field int staticValue>
4 aload_0
5 iload_2
6 putfield #2 <Field int instanceValue>
9 return
```

A6. The sixth excerpt is the same example application in 5 after modification has been performed. The modifications are highlighted in **bold**.

```
Method void setValues(int, int)
0 iload_1
1 putstatic #3 <Field int staticValue>
4 ldc #4 <String "example">
6 iconst_0
7 invokestatic #5 <Method void alert(java.lang.Object, int)>
10 aload_0
11 iload_2
12 putfield #2 <Field int instanceValue>
15 aload_0
16 iconst_1
17 invokestatic #5 <Method void alert(java.lang.Object, int)>
20 return
```

A7. The seventh excerpt is the source-code of the example application used in excerpt 5 and 6.

```
import java.lang.*;

public class example{

    /** Shared static field. */
    public static int staticValue = 0;

    /** Shared instance field. */
    public int instanceValue = 0;

    /** Example method that writes to memory (instance field). */
    public void setValues(int a, int b){

        staticValue = a;

        instanceValue = b;

    }

}
```

A8. The eighth excerpt is the source-code of FieldAlert, which alerts the "distributed run-time" to propagate a changed value.

```

import java.lang.*;
import java.util.*;
import java.net.*;
import java.io.*;

public class FieldAlert{

    /** Table of alerts. */
    public final static Hashtable alerts = new Hashtable();

    /** Object handle. */
    public Object reference = null;

    /** Table of field alerts for this object. */
    public boolean[] fieldAlerts = null;

    /** Constructor. */
    public FieldAlert(Object o, int initialFieldCount){
        reference = o;
        fieldAlerts = new boolean[initialFieldCount];
    }

    /** Called when an application modifies a value. (Both objects and
    classes) */
    public static void alert(Object o, int fieldID){

        // Lock the alerts table.
        synchronized (alerts){

            FieldAlert alert = (FieldAlert) alerts.get(o);

            if (alert == null){ // This object hasn't been alerted already,
                               // so add to alerts table.
                alert = new FieldAlert(o, fieldID + 1);
                alerts.put(o, alert);
            }

            if (fieldID >= alert.fieldAlerts.length){
                // Ok, enlarge fieldAlerts array.
                boolean[] b = new boolean[fieldID+1];
                System.arraycopy(alert.fieldAlerts, 0, b, 0,
                                alert.fieldAlerts.length);
                alert.fieldAlerts = b;
            }

            // Record the alert.
            alert.fieldAlerts[fieldID] = true;

            // Mark as pending.
            FieldSend.pending = true; // Signal that there is one or more
                                     // propagations waiting.

            // Finally, notify the waiting FieldSend thread(s)
            if (FieldSend.waiting){
                FieldSend.waiting = false;
                alerts.notify();
            }
        }
    }
}

```

A9. The ninth excerpt is the source-code of FieldSend, which propagates changes values alerted to it via FieldAlert.

```
import java.lang.*;
import java.lang.reflect.*;
import java.util.*;
import java.net.*;
import java.io.*;

public class FieldSend implements Runnable{

    /** Protocol specific values. */
    public final static int CLOSE = -1;
    public final static int NACK = 0;
    public final static int ACK = 1;
    public final static int PROPAGATE_OBJECT = 10;
    public final static int PROPAGATE_CLASS = 20;

    /** FieldAlert network values. */
    public final static String group =
        System.getProperty("FieldAlert_network_group");
    public final static int port =
        Integer.parseInt(System.getProperty("FieldAlert_network_port"));

    /** Table of global ID's for local objects. (hashcode-to-globalID
    mappings) */
    public final static Hashtable objectToGlobalID = new Hashtable();

    /** Table of global ID's for local classnames. (classname-to-globalID
    mappings) */
    public final static Hashtable classNameToGlobalID = new Hashtable();

    /** Pending. True if a propagation is pending. */
    public static boolean pending = false;

    /** Waiting. True if the FieldSend thread(s) are waiting. */
    public static boolean waiting = false;

    /** Background send thread. Propagates values as this thread is alerted
    to their alteration. */
    public void run(){

        System.out.println("FieldAlert_network_group=" + group);
        System.out.println("FieldAlert_network_port=" + port);

        try{

            // Create a DatagramSocket to send propagated field values.
            DatagramSocket datagramSocket =
                new DatagramSocket(port, InetAddress.getByName(group));

            // Next, create the buffer and packet for all transmissions.
            byte[] buffer = new byte[512];           // Working limit of 512 bytes
                                                    // per packet.

            DatagramPacket datagramPacket =
                new DatagramPacket(buffer, 0, buffer.length);

            while (!Thread.interrupted()){

                Object[] entries = null;

                // Lock the alerts table.
                synchronized (FieldAlert.alerts){
```



```

// Await for an alert to propagate something.
while (!pending){
    waiting = true;
    FieldAlert.alerts.wait();
    waiting = false;
}

pending = false;

entries = FieldAlert.alerts.entrySet().toArray();

// Clear alerts once we have copied them.
FieldAlert.alerts.clear();
}

// Process each object alert in turn.
for (int i=0; i<entries.length; i++){

    FieldAlert alert = (FieldAlert) entries[i];

    int index = 0;
    datagramPacket.setLength(buffer.length);

    Object reference = null;
    if (alert.reference instanceof String){
        // PROPAGATE_CLASS field operation.

        buffer[index++] = (byte) ((PROPAGATE_CLASS >> 24) & 0xff);
        buffer[index++] = (byte) ((PROPAGATE_CLASS >> 16) & 0xff);
        buffer[index++] = (byte) ((PROPAGATE_CLASS >> 8) & 0xff);
        buffer[index++] = (byte) ((PROPAGATE_CLASS >> 0) & 0xff);

        String name = (String) alert.reference;
        int length = name.length();
        buffer[index++] = (byte) ((length >> 24) & 0xff);
        buffer[index++] = (byte) ((length >> 16) & 0xff);
        buffer[index++] = (byte) ((length >> 8) & 0xff);
        buffer[index++] = (byte) ((length >> 0) & 0xff);

        byte[] bytes = name.getBytes();
        System.arraycopy(bytes, 0, buffer, index, length);
        index += length;
    }else{
        // PROPAGATE_OBJECT field operation.

        buffer[index++] =
            (byte) ((PROPAGATE_OBJECT >> 24) & 0xff);
        buffer[index++] =
            (byte) ((PROPAGATE_OBJECT >> 16) & 0xff);
        buffer[index++] = (byte) ((PROPAGATE_OBJECT >> 8) & 0xff);
        buffer[index++] = (byte) ((PROPAGATE_OBJECT >> 0) & 0xff);

        int globalID = ((Integer)
            objectToGlobalID.get(alert.reference)).intValue();

        buffer[index++] = (byte) ((globalID >> 24) & 0xff);
        buffer[index++] = (byte) ((globalID >> 16) & 0xff);
        buffer[index++] = (byte) ((globalID >> 8) & 0xff);
        buffer[index++] = (byte) ((globalID >> 0) & 0xff);

        reference = alert.reference;
    }

    // Use reflection to get a table of fields that correspond to
    // the field indexes used internally.

```

```

Field[] fields = null;
if (reference == null){
    fields = FieldLoader.loadClass((String)
        alert.reference).getDeclaredFields();
}else{
    fields = alert.reference.getClass().getDeclaredFields();
}

// Now encode in batch mode the fieldID/value pairs.
for (int j=0; j<alert.fieldAlerts.length; j++){
    if (alert.fieldAlerts[j] == false)
        continue;

    buffer[index++] = (byte) ((j >> 24) & 0xff);
    buffer[index++] = (byte) ((j >> 16) & 0xff);
    buffer[index++] = (byte) ((j >> 8) & 0xff);
    buffer[index++] = (byte) ((j >> 0) & 0xff);

    // Encode value.
    Class type = fields[j].getType();
    if (type == Boolean.TYPE){
        buffer[index++] = (byte)
            (fields[j].getBoolean(reference)? 1 : 0);
    }else if (type == Byte.TYPE){
        buffer[index++] = fields[j].getByte(reference);
    }else if (type == Short.TYPE){
        short v = fields[j].getShort(reference);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else if (type == Character.TYPE){
        char v = fields[j].getChar(reference);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else if (type == Integer.TYPE){
        int v = fields[j].getInt(reference);
        buffer[index++] = (byte) ((v >> 24) & 0xff);
        buffer[index++] = (byte) ((v >> 16) & 0xff);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else if (type == Float.TYPE){
        int v = Float.floatToIntBits(
            fields[j].getFloat(reference));
        buffer[index++] = (byte) ((v >> 24) & 0xff);
        buffer[index++] = (byte) ((v >> 16) & 0xff);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else if (type == Long.TYPE){
        long v = fields[j].getLong(reference);
        buffer[index++] = (byte) ((v >> 56) & 0xff);
        buffer[index++] = (byte) ((v >> 48) & 0xff);
        buffer[index++] = (byte) ((v >> 40) & 0xff);
        buffer[index++] = (byte) ((v >> 32) & 0xff);
        buffer[index++] = (byte) ((v >> 24) & 0xff);
        buffer[index++] = (byte) ((v >> 16) & 0xff);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else if (type == Double.TYPE){
        long v = Double.doubleToLongBits(
            fields[j].getDouble(reference));
        buffer[index++] = (byte) ((v >> 56) & 0xff);
        buffer[index++] = (byte) ((v >> 48) & 0xff);
        buffer[index++] = (byte) ((v >> 40) & 0xff);
        buffer[index++] = (byte) ((v >> 32) & 0xff);
        buffer[index++] = (byte) ((v >> 24) & 0xff);
        buffer[index++] = (byte) ((v >> 16) & 0xff);
        buffer[index++] = (byte) ((v >> 8) & 0xff);
        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }
}

```

```

        buffer[index++] = (byte) ((v >> 0) & 0xff);
    }else{
        throw new AssertionError("Unsupported type.
    }

}

// Now set the length of the datagramPacket.
datagramPacket.setLength(index);

// Now send the packet.
datagramSocket.send(datagramPacket);

}

}

} catch (Exception e){
    throw new AssertionError("Exception: " + e.toString());
}

}

```

A10. The tenth excerpt is the source-code of FieldReceive, which receives propagated changed values sent via FieldSend.

```
import java.lang.*;
import java.lang.reflect.*;
import java.util.*;
import java.net.*;
import java.io.*;

public class FieldReceive implements Runnable{

    /** Protocol specific values. */
    public final static int CLOSE = -1;
    public final static int NACK = 0;
    public final static int ACK = 1;
    public final static int PROPAGATE_OBJECT = 10;
    public final static int PROPAGATE_CLASS = 20;

    /** FieldAlert network values. */
    public final static String group =
        System.getProperty("FieldAlert_network_group");
    public final static int port =
        Integer.parseInt(System.getProperty("FieldAlert_network_port"));

    /** Table of global ID's for local objects. (globalID-to-hashcode
        mappings) */
    public final static Hashtable globalIDToObject = new Hashtable();

    /** Table of global ID's for local classnames. (globalID-to-classname
        mappings) */
    public final static Hashtable globalIDToClassName = new Hashtable();

    /** Called when an application is to acquire a lock. */
    public void run(){

        System.out.println("FieldAlert_network_group=" + group);
        System.out.println("FieldAlert_network_port=" + port);

        try{
```

```

// Create a DatagramSocket to send propagated field values from
MulticastSocket multicastSocket = new MulticastSocket(port);
multicastSocket.joinGroup(InetAddress.getByName(group));

// Next, create the buffer and packet for all transmissions.
byte[] buffer = new byte[512];           // Working limit of 512
                                           // bytes per packet.
DatagramPacket datagramPacket =
    new DatagramPacket(buffer, 0, buffer.length);

while (!Thread.interrupted()){

    // Make sure to reset length.
    datagramPacket.setLength(buffer.length);

    // Receive the next available packet.
    multicastSocket.receive(datagramPacket);

    int index = 0, length = datagramPacket.getLength();

    // Decode the command.
    int command = (int) (((buffer[index++] & 0xff) << 24)
        | ((buffer[index++] & 0xff) << 16)
        | ((buffer[index++] & 0xff) << 8)
        | (buffer[index++] & 0xff));

    if (command == PROPAGATE_OBJECT){ // Propagate operation for
                                      // object fields.

        // Decode global id.
        int globalID = (int) (((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));

        // Now, need to resolve the object in question.
        Object reference = globalIDToObject.get(
            new Integer(globalID));

        // Next, get the array of fields for this object.
        Field[] fields = reference.getClass().getDeclaredFields();

        while (index < length){

            // Decode the field id.
            int fieldID = (int) (((buffer[index++] & 0xff) << 24)
                | ((buffer[index++] & 0xff) << 16)
                | ((buffer[index++] & 0xff) << 8)
                | (buffer[index++] & 0xff));

            // Determine value length based on corresponding field
            // type.
            Field field = fields[fieldID];
            Class type = field.getType();
            if (type == Boolean.TYPE){
                boolean v = (buffer[index++] == 1 ? true : false);
                field.setBoolean(reference, v);
            }else if (type == Byte.TYPE){
                byte v = buffer[index++];
                field.setByte(reference, v);
            }else if (type == Short.TYPE){
                short v = (short) (((buffer[index++] & 0xff) << 8)
                    | (buffer[index++] & 0xff));
                field.setShort(reference, v);
            }else if (type == Character.TYPE){
                char v = (char) (((buffer[index++] & 0xff) << 8)

```

```

        | (buffer[index++] & 0xff));
        field.setChar(reference, v);
    }else if (type == Integer.TYPE){
        int v = (int) (((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));
        field.setInt(reference, v);
    }else if (type == Float.TYPE){
        int v = (int) (((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));
        field.setFloat(reference, Float.intBitsToFloat(v));
    }else if (type == Long.TYPE){
        long v = (long) (((buffer[index++] & 0xff) << 56)
            | ((buffer[index++] & 0xff) << 48)
            | ((buffer[index++] & 0xff) << 40)
            | ((buffer[index++] & 0xff) << 32)
            | ((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));
        field.setLong(reference, v);
    }else if (type == Double.TYPE){
        long v = (long) (((buffer[index++] & 0xff) << 56)
            | ((buffer[index++] & 0xff) << 48)
            | ((buffer[index++] & 0xff) << 40)
            | ((buffer[index++] & 0xff) << 32)
            | ((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));
        field.setDouble(reference, Double.longBitsToDouble(v));
    }else{
        throw new AssertionError("Unsupported type.");
    }
}

}

}else if (command == PROPAGATE_CLASS){ // Propagate an update
    // to class fields.

    // Decode the classname.
    int nameLength = (int) (((buffer[index++] & 0xff) << 24)
        | ((buffer[index++] & 0xff) << 16)
        | ((buffer[index++] & 0xff) << 8)
        | (buffer[index++] & 0xff));
    String name = new String(buffer, index, nameLength);
    index += nameLength;

    // Next, get the array of fields for this class.
    Field[] fields =
        FieldLoader.loadClass(name).getDeclaredFields();

    // Decode all batched fields included in this propagation
    // packet.
    while (index < length){

        // Decode the field id.
        int fieldID = (int) (((buffer[index++] & 0xff) << 24)
            | ((buffer[index++] & 0xff) << 16)
            | ((buffer[index++] & 0xff) << 8)
            | (buffer[index++] & 0xff));

        // Determine field type to determine value length.
        Field field = fields[fieldID];

```

[illegible]

This excerpt is the source-code of FieldLoader, which modifies an application as it is being loaded.

```
import java.lang.*;
import java.io.*;
import java.net.*;

public class FieldLoader extends URLClassLoader{

    public FieldLoader(URL[] urls){
        super(urls);
    }

    protected Class findClass(String name)
    throws ClassNotFoundException{

        ClassFile cf = null;

        try{

            BufferedInputStream in =
                new BufferedInputStream(findResource{
                    name.replace('.', '/').concat(".class").openStream());

            cf = new ClassFile(in);

        }catch (Exception e){throw new ClassNotFoundException(e.toString());}

        // Class-wide pointers to the ldc and alert index.
        int ldcindex = -1;
        int alertindex = -1;

        for (int i=0; i<cf.methods_count; i++){
            for (int j=0; j<cf.methods[i].attributes_count; j++){
                if (!(cf.methods[i].attributes[j] instanceof Code_attribute))
                    continue;

                Code_attribute ca = (Code_attribute)
                    cf.methods[i].attributes[j];

                boolean changed = false;

                for (int z=0; z<ca.code.length; z++){
                    if ((ca.code[z][0] & 0xff) == 179){ // Opcode for a
PUTSTATIC                                     // instruction.

                        changed = true;

                        // The code below only supports fields in this class.
                        // Thus, first off, check that this field is local to this
                        // class.
                        CONSTANT_Fieldref_info fi = (CONSTANT_Fieldref_info)
                            cf.constant_pool[(int) ((ca.code[z][1] & 0xff) << 8) |
                            (ca.code[z][2] & 0xff)]];
                        CONSTANT_Class_info ci = (CONSTANT_Class_info)
                            cf.constant_pool[fi.class_index];
                        String className =
                            cf.constant_pool[ci.name_index].toString();
                        if (!name.equals(className)){
                            throw new AssertionError("This code only supports
fields "
                                "local to this class");
                        }

                        // Ok, now search for the fields name and index.
```

```

int index = 0;
CONSTANT_NameAndType_info ni = (CONSTANT_NameAndType_info)
    cf.constant_pool[fi.name_and_type_index];
String fieldName =
    cf.constant_pool[ni.name_index].toString();
for (int a=0; a<cf.fields_count; a++){
    String fn = cf.constant_pool[
        cf.fields[a].name_index].toString();
    if (fieldName.equals(fn)){
        index = a;
        break;
    }
}
// Next, realign the code array, making room for the
// insertions.
byte[] code2 = new byte[ca.code.length+3];
System.arraycopy(ca.code, 0, code2, 0, z+1);
System.arraycopy(ca.code, z+1, code2, z+4,
    ca.code.length-(z+1));
ca.code = code2;

// Next, insert the LDC_W instruction.
if (ldcindex == -1){
    CONSTANT_String_info csi =
        new CONSTANT_String_info(ci.name_index);
    cp_info[] cpi = new cp_info[cf.constant_pool.length+1];
    System.arraycopy(cf.constant_pool, 0, cpi, 0,
        cf.constant_pool.length);
    cpi[cpi.length - 1] = csi;
    ldcindex = cpi.length-1;
    cf.constant_pool = cpi;
    cf.constant_pool_count++;
}
ca.code[z+1] = new byte[3];
ca.code[z+1][0] = (byte) 19;
ca.code[z+1][1] = (byte) ((ldcindex >> 8) & 0xff);
ca.code[z+1][2] = (byte) (ldcindex & 0xff);

// Next, insert the SIPUSH instruction.
ca.code[z+2] = new byte[3];
ca.code[z+2][0] = (byte) 17;
ca.code[z+2][1] = (byte) ((index >> 8) & 0xff);
ca.code[z+2][2] = (byte) (index & 0xff);

// Finally, insert the INVOKESTATIC instruction.
if (alertindex == -1){
    // This is the first time this class is encountering
    // alert instruction, so have to add it to the constant
    // pool.
    cp_info[] cpi = new cp_info[cf.constant_pool.length+6];
    System.arraycopy(cf.constant_pool, 0, cpi, 0,
        cf.constant_pool.length);
    cf.constant_pool = cpi;
    cf.constant_pool_count += 6;

    CONSTANT_Utf8_info u1 =
        new CONSTANT_Utf8_info("FieldAlert");
    cf.constant_pool[cf.constant_pool.length-6] = u1;

    CONSTANT_Class_info c1 = new CONSTANT_Class_info(
        cf.constant_pool_count-6);
    cf.constant_pool[cf.constant_pool.length-5] = c1;

    u1 = new CONSTANT_Utf8_info("alert");
    cf.constant_pool[cf.constant_pool.length-4] = u1;
}

```

the


```

        u1 = new CONSTANT_Utf8_info("(Ljava/lang/Object;I)V");
        cf.constant_pool[cf.constant_pool.length-3] = u1;

        CONSTANT_NameAndType_info n1 =
            new CONSTANT_NameAndType_info(
                cf.constant_pool.length-4, cf.constant_pool.length-
3);
        cf.constant_pool[cf.constant_pool.length-2] = n1;

        CONSTANT_Methodref_info m1 = new
CONSTANT_Methodref_info(
2);
            cf.constant_pool.length-5, cf.constant_pool.length-
            cf.constant_pool[cf.constant_pool.length-1] = m1;
            alertindex = cf.constant_pool.length-1;
        }
        ca.code[z+3] = new byte[3];
        ca.code[z+3][0] = (byte) 184;
        ca.code[z+3][1] = (byte) ((alertindex >> 8) & 0xff);
        ca.code[z+3][2] = (byte) (alertindex & 0xff);

        // And lastly, increase the CODE_LENGTH and
ATTRIBUTE_LENGTH
        // values.
        ca.code_length += 9;
        ca.attribute_length += 9;

    }

}

// If we changed this method, then increase the stack size by
one.
if (changed){
    ca.max_stack++;
}

}

}

try{
    ByteArrayOutputStream out = new ByteArrayOutputStream();
    cf.serialize(out);

    byte[] b = out.toByteArray();

    return defineClass(name, b, 0, b.length);

}catch (Exception e){
    throw new ClassNotFoundException(name);
}

}

}

```

A12. Attribute_info.java

Convenience class for representing attribute_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** This abstract class represents all types of attribute_info
 * that are used in the JVM specifications.

```

```

*
* All new attribute_info subclasses are to always inherit from this
* class.
*/
public abstract class attribute_info{

    public int attribute_name_index;
    public int attribute_length;

    /** This is used by subclasses to register themselves
     * to their parent classFile.
     */
    attribute_info(ClassFile cf){}

    /** Used during input serialization by ClassFile only. */
    attribute_info(ClassFile cf, DataInputStream in)
        throws IOException{
        attribute_name_index = in.readChar();
        attribute_length = in.readInt();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeChar(attribute_name_index);
        out.writeInt(attribute_length);
    }

    /** This class represents an unknown attribute_info that
     * this current version of classfile specification does
     * not understand.
     */
    public final static class Unknown extends attribute_info{

        byte[] info;

        /** Used during input serialization by ClassFile only. */
        Unknown(ClassFile cf, DataInputStream in)
            throws IOException{
            super(cf, in);
            info = new byte[attribute_length];
            in.read(info, 0, attribute_length);
        }

        /** Used during output serialization by ClassFile only. */
        void serialize(DataOutputStream out)
            throws IOException{
            ByteArrayOutputStream baos = new ByteArrayOutputStream();

            super.serialize(out);
            out.write(info, 0, attribute_length);
        }
    }
}

```

A13. ClassFile.java

Convenience class for representing ClassFile structures.

```

import java.lang.*;
import java.io.*;
import java.util.*;

/** The ClassFile follows verbatim from the JVM specification. */
public final class ClassFile {

    public int magic;
    public int minor_version;

```

```

public int major_version;
public int constant_pool_count;
public cp_info[] constant_pool;
public int access_flags;
public int this_class;
public int super_class;
public int interfaces_count;
public int[] interfaces;
public int fields_count;
public field_info[] fields;
public int methods_count;
public method_info[] methods;
public int attributes_count;
public attribute_info[] attributes;

/** Constructor. Takes in a byte stream representation and transforms
 * each of the attributes in the ClassFile into objects to allow for
 * easier manipulation.
 */
public ClassFile(InputStream ins)
    throws IOException{
    DataInputStream in = (ins instanceof DataInputStream ?
        (DataInputStream) ins : new DataInputStream(ins));
    magic = in.readInt();
    minor_version = in.readChar();
    major_version = in.readChar();
    constant_pool_count = in.readChar();
    constant_pool = new cp_info[constant_pool_count];
    for (int i=1; i<constant_pool_count; i++){
        in.mark(1);
        int s = in.read();
        in.reset();
        switch (s){
            case 1:
                constant_pool[i] = new CONSTANT_Utf8_info(this, in);
                break;
            case 3:
                constant_pool[i] = new CONSTANT_Integer_info(this, in);
                break;
            case 4:
                constant_pool[i] = new CONSTANT_Float_info(this, in);
                break;
            case 5:
                constant_pool[i] = new CONSTANT_Long_info(this, in);
                i++;
                break;
            case 6:
                constant_pool[i] = new CONSTANT_Double_info(this, in);
                i++;
                break;
            case 7:
                constant_pool[i] = new CONSTANT_Class_info(this, in);
                break;
            case 8:
                constant_pool[i] = new CONSTANT_String_info(this, in);
                break;
            case 9:
                constant_pool[i] = new CONSTANT_Fieldref_info(this, in);
                break;
            case 10:
                constant_pool[i] = new CONSTANT_Methodref_info(this,
in);
                break;
            case 11:
                constant_pool[i] =
                    new CONSTANT_InterfaceMethodref_info(this, in);
                break;

```

```

        case 12:
            constant_pool[i] = new CONSTANT_NameAndType_info(this,
in);
            break;
        default:
            throw new ClassFormatError("Invalid ConstantPoolTag");
    }
}

access_flags = in.readChar();
this_class = in.readChar();
super_class = in.readChar();
interfaces_count = in.readChar();
interfaces = new int[interfaces_count];
for (int i=0; i<interfaces_count; i++)
    interfaces[i] = in.readChar();
fields_count = in.readChar();
fields = new field_info[fields_count];
for (int i=0; i<fields_count; i++) {
    fields[i] = new field_info(this, in);
}
methods_count = in.readChar();
methods = new method_info[methods_count];
for (int i=0; i<methods_count; i++) {
    methods[i] = new method_info(this, in);
}
attributes_count = in.readChar();
attributes = new attribute_info[attributes_count];
for (int i=0; i<attributes_count; i++){
    in.mark(2);
    String s = constant_pool[in.readChar()].toString();
    in.reset();
    if (s.equals("SourceFile"))
        attributes[i] = new SourceFile_attribute(this, in);
    else if (s.equals("Deprecated"))
        attributes[i] = new Deprecated_attribute(this, in);
    else if (s.equals("InnerClasses"))
        attributes[i] = new InnerClasses_attribute(this, in);
    else
        attributes[i] = new attribute_info.Unknown(this, in);
}
}

/** Serializes the ClassFile object into a byte stream. */
public void serialize(OutputStream o)
    throws IOException{
    DataOutputStream out = (o instanceof DataOutputStream ?
        (DataOutputStream) o : new DataOutputStream(o));
    out.writeInt(magic);
    out.writeChar(minor_version);
    out.writeChar(major_version);
    out.writeChar(constant_pool_count);
    for (int i=1; i<constant_pool_count; i++){
        constant_pool[i].serialize(out);
        if (constant_pool[i] instanceof CONSTANT_Long_info ||
            constant_pool[i] instanceof CONSTANT_Double_info)
            i++;
    }
    out.writeChar(access_flags);
    out.writeChar(this_class);
    out.writeChar(super_class);
    out.writeChar(interfaces_count);
    for (int i=0; i<interfaces_count; i++)
        out.writeChar(interfaces[i]);
    out.writeChar(fields_count);
    for (int i=0; i<fields_count; i++)
        fields[i].serialize(out);
    out.writeChar(methods_count);

```

```

        for (int i=0; i<methods_count; i++)
            methods[i].serialize(out);
        out.writeChar(attributes_count);
        for (int i=0; i<attributes_count; i++)
            attributes[i].serialize(out);
        // Flush the outputstream just to make sure.
        out.flush();
    }
}

```

A14. Code_attribute.java

Convenience class for representing Code_attribute structures within ClassFiles.

```

import java.util.*;
import java.lang.*;
import java.io.*;

/**
 * The code[] is stored as a 2D array. */
public final class Code_attribute extends attribute_info{

    public int max_stack;
    public int max_locals;
    public int code_length;
    public byte[][] code;
    public int exception_table_length;
    public exception_table[] exception_table;
    public int attributes_count;
    public attribute_info[] attributes;

    /** Internal class that handles the exception table. */
    public final static class exception_table{
        public int start_pc;
        public int end_pc;
        public int handler_pc;
        public int catch_type;
    }

    /** Constructor called only by method_info. */
    Code_attribute(ClassFile cf, int ani, int al, int ms, int ml, int cl,
        byte[][] cd, int etl, exception_table[] et, int ac,
        attribute_info[] a){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        max_stack = ms;
        max_locals = ml;
        code_length = cl;
        code = cd;
        exception_table_length = etl;
        exception_table = et;
        attributes_count = ac;
        attributes = a;
    }

    /** Used during input serialization by ClassFile only. */
    Code_attribute(ClassFile cf, DataInputStream in)
        Throws IOException{
        super(cf, in);
        max_stack = in.readChar();
        max_locals = in.readChar();
        code_length = in.readInt();
        code = new byte[code_length][];

        int i = 0;
        for (int pos=0; pos<code_length; i++){

```

```
in.mark(1);
int s = in.read();
in.reset();
switch (s){
    case 16:
    case 18:
    case 21:
    case 22:
    case 23:
    case 24:
    case 25:
    case 54:
    case 55:
    case 56:
    case 57:
    case 58:
    case 169:
    case 188:
    case 196:
        code[i] = new byte[2];
        break;
    case 17:
    case 19:
    case 20:
    case 132:
    case 153:
    case 154:
    case 155:
    case 156:
    case 157:
    case 158:
    case 159:
    case 160:
    case 161:
    case 162:
    case 163:
    case 164:
    case 165:
    case 166:
    case 167:
    case 168:
    case 178:
    case 179:
    case 180:
    case 181:
    case 182:
    case 183:
    case 184:
    case 187:
    case 189:
    case 192:
    case 193:
    case 198:
    case 199:
    case 209:
        code[i] = new byte[3];
        break;
    case 197:
        code[i] = new byte[4];
        break;
    case 185:
    case 200:
    case 201:
        code[i] = new byte[5];
        break;
    case 170:{
        int pad = 3 - (pos % 4);
```

```

        in.mark(pad+13);                // highbyte
        in.skipBytes(pad+5);           // lowbyte
        int low = in.readInt();
        code[i] =
            new byte[pad + 13 + ((in.readInt() - low + 1) * 4)];
        in.reset();
        break;
    )case 171:
        int pad = 3 - (pos % 4);
        in.mark(pad+9);
        in.skipBytes(pad+5);
        code[i] = new byte[pad + 9 + (in.readInt() * 8)];
        in.reset();
        break;
    )default:
        code[i] = new byte[1];
    )
    in.read(code[i], 0, code[i].length);
    pos += code[i].length;
}

// adjust the array to the new size and store the size
byte[i][] temp = new byte[i][];
System.arraycopy(code, 0, temp, 0, i);
code = temp;

exception_table_length = in.readChar();
exception_table =
    new Code_attribute.exception_table[exception_table_length];
for (i=0; i<exception_table_length; i++){
    exception_table[i] = new exception_table();
    exception_table[i].start_pc = in.readChar();
    exception_table[i].end_pc = in.readChar();
    exception_table[i].handler_pc = in.readChar();
    exception_table[i].catch_type = in.readChar();
}
attributes_count = in.readChar();
attributes = new attribute_info[attributes_count];
for (i=0; i<attributes_count; i++){
    in.mark(2);
    String s = cf.constant_pool[in.readChar()].toString();
    in.reset();
    if (s.equals("LineNumberTable"))
        attributes[i] = new LineNumberTable_attribute(cf, in);
    else if (s.equals("LocalVariableTable"))
        attributes[i] = new LocalVariableTable_attribute(cf, in);
    else
        attributes[i] = new attribute_info.Unknown(cf, in);
}
}

/** Used during output serialization by ClassFile only.
 */
void serialize(DataOutputStream out)
    throws IOException{
    attribute_length = 12 + code_length +
        (exception_table_length * 8);
    for (int i=0; i<attributes_count; i++){
        attribute_length += attributes[i].attribute_length + 6;
    }
    super.serialize(out);
    out.writeChar(max_stack);
    out.writeChar(max_locals);
    out.writeInt(code_length);
    for (int i=0, pos=0; pos<code_length; i++){
        out.write(code[i], 0, code[i].length);
        pos += code[i].length;
    }
}

```

```

        out.writeChar(exception_table_length);
        for (int i=0; i<exception_table_length; i++){
            out.writeChar(exception_table[i].start_pc);
            out.writeChar(exception_table[i].end_pc);
            out.writeChar(exception_table[i].handler_pc);
            out.writeChar(exception_table[i].catch_type);
        }
        out.writeChar(attributes_count);
        for (int i=0; i<attributes_count; i++)
            attributes[i].serialize(out);
    }
}

```

A15. CONSTANT_Class_info.java

Convenience class for representing CONSTANT_Class_info structures within

ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Class subtype of a constant pool entry. */
public final class CONSTANT_Class_info extends cp_info{

    /** The index to the name of this class. */
    public int name_index = 0;

    /** Convenience constructor. */
    public CONSTANT_Class_info(int index) {
        tag = 7;
        name_index = index;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Class_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 7)
            throw new ClassFormatError();
        name_index = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(name_index);
    }
}

```

A16. CONSTANT_Double_info.java

Convenience class for representing CONSTANT_Double_info structures within

ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Double subtype of a constant pool entry. */
public final class CONSTANT_Double_info extends cp_info{

```



```

    /** The actual value. */
    public double bytes;

    public CONSTANT_Double_info(double d){
        tag = 6;
        bytes = d;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Double_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 6)
            throw new ClassFormatError();
        bytes = in.readDouble();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeDouble(bytes);
        long l = Double.doubleToLongBits(bytes);
    }
}

```

A17. CONSTANT_Fieldref_info.java

Convenience class for representing CONSTANT_Fieldref_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Fieldref subtype of a constant pool entry. */
public final class CONSTANT_Fieldref_info extends cp_info{

    /** The index to the class that this field is referencing to. */
    public int class_index;

    /** The name and type index this field is referencing to. */
    public int name_and_type_index;

    /** Convenience constructor. */
    public CONSTANT_Fieldref_info(int class_index, int name_and_type_index)
    {
        tag = 9;
        this.class_index = class_index;
        this.name_and_type_index = name_and_type_index;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Fieldref_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 9)
            throw new ClassFormatError();
        class_index = in.readChar();
        name_and_type_index = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(class_index);
    }
}

```

```

        out.writeChar(name_and_type_index);
    }
}

```

A18. CONSTANT_Float_info.java

Convenience class for representing CONSTANT_Float_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Float subtype of a constant pool entry. */
public final class CONSTANT_Float_info extends cp_info{

    /** The actual value. */
    public float bytes;

    public CONSTANT_Float_info(float f){
        tag = 4;
        bytes = f;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Float_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 4)
            throw new ClassFormatError();
        bytes = in.readFloat();
    }

    /** Used during output serialization by ClassFile only. */
    public void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(4);
        out.writeFloat(bytes);
    }
}

```

A19. CONSTANT_Integer_info.java

Convenience class for representing CONSTANT_Integer_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Integer subtype of a constant pool entry. */
public final class CONSTANT_Integer_info extends cp_info{

    /** The actual value. */
    public int bytes;

    public CONSTANT_Integer_info(int b) {
        tag = 3;
        bytes = b;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Integer_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);

```

```

        if (tag != 3)
            throw new ClassFormatException();
        bytes = in.readInt();
    }

    /** Used during output serialization by ClassFile only. */
    public void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeInt(bytes);
    }

}

A20. CONSTANT_InterfaceMethodref_info.java
Convenience class for representing CONSTANT_InterfaceMethodref_info structures
within ClassFiles.

import java.lang.*;
import java.io.*;

/** InterfaceMethodref subtype of a constant pool entry.
 */
public final class CONSTANT_InterfaceMethodref_info extends cp_info{

    /** The index to the class that this field is referencing to. */
    public int class_index;

    /** The name and type index this field is referencing to. */
    public int name_and_type_index;

    public CONSTANT_InterfaceMethodref_info(int class_index,
                                           int name_and_type_index) {
        tag = 11;
        this.class_index = class_index;
        this.name_and_type_index = name_and_type_index;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_InterfaceMethodref_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 11)
            throw new ClassFormatException();
        class_index = in.readChar();
        name_and_type_index = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(class_index);
        out.writeChar(name_and_type_index);
    }

}

A21. CONSTANT_Long_info.java
Convenience class for representing CONSTANT_Long_info structures within
ClassFiles.

import java.lang.*;

```

```

import java.io.*;

/** Long subtype of a constant pool entry. */
public final class CONSTANT_Long_info extends cp_info{

    /** The actual value. */
    public long bytes;

    public CONSTANT_Long_info(long b){
        tag = 5;
        bytes = b;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Long_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 5)
            throw new ClassFormatException();
        bytes = in.readLong();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeLong(bytes);
    }
}

```

A22. CONSTANT_Methodref_info.java

Convenience class for representing CONSTANT_Methodref_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Methodref subtype of a constant pool entry.
 */
public final class CONSTANT_Methodref_info extends cp_info{

    /** The index to the class that this field is referencing to. */
    public int class_index;

    /** The name and type index this field if referencing to. */
    public int name_and_type_index;

    public CONSTANT_Methodref_info(int class_index, int name_and_type_index)
    {
        tag = 10;
        this.class_index = class_index;
        this.name_and_type_index = name_and_type_index;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_Methodref_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 10)
            throw new ClassFormatException();
        class_index = in.readChar();
        name_and_type_index = in.readChar();
    }
}

```

```

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(class_index);
        out.writeChar(name_and_type_index);
    }
}

```

A23. CONSTANT_NameAndType_info.java

Convenience class for representing CONSTANT_NameAndType_info structures within ClassFiles.

```

import java.io.*;
import java.lang.*;

/** NameAndType subtype of a constant pool entry.
 */
public final class CONSTANT_NameAndType_info extends cp_info{

    /** The index to the Utf8 that contains the name. */
    public int name_index;

    /** The index fo the Utf8 that constains the signature. */
    public int descriptor_index;

    public CONSTANT_NameAndType_info(int name_index, int descriptor_index) {
        tag = 12;
        this.name_index = name_index;
        this.descriptor_index = descriptor_index;
    }

    /** Used during input serialization by ClassFile only. */
    CONSTANT_NameAndType_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 12){
            throw new ClassFormatException();
        }
        name_index = in.readChar();
        descriptor_index = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(name_index);
        out.writeChar(descriptor_index);
    }
}

```

A24. CONSTANT_String_info.java

Convenience class for representing CONSTANT_String_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

```

```

/** String subtype of a constant pool entry.
 */
public final class CONSTANT_String_info extends cp_info{

    /** The index to the actual value of the string. */
    public int string_index;

    public CONSTANT_String_info(int value) {
        tag = 8;
        string_index = value;
    }

    /** ONLY TO BE USED BY CLASSFILE! */
    public CONSTANT_String_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 8)
            throw new ClassFormatException();
        string_index = in.readChar();
    }

    /** Output serialization, ONLY TO BE USED BY CLASSFILE! */
    public void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(string_index);
    }

}

```

A25. CONSTANT_Utf8_info.java

Convenience class for representing CONSTANT_Utf8_info structures within ClassFiles.

```

import java.io.*;
import java.lang.*;

/** Utf8 subtype of a constant pool entry.
 * We internally represent the Utf8 info byte array
 * as a String.
 */
public final class CONSTANT_Utf8_info extends cp_info{

    /** Length of the byte array. */
    public int length;

    /** The actual bytes, represented by a String. */
    public String bytes;

    /** This constructor should be used for the purpose
     * of part creation. It does not set the parent
     * ClassFile reference.
     */
    public CONSTANT_Utf8_info(String s) {
        tag = 1;
        length = s.length();
        bytes = s;
    }

    /** Used during input serialization by ClassFile only. */
    public CONSTANT_Utf8_info(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        if (tag != 1)

```

```

        throw new ClassFormatError();
        length = in.readChar();
        byte[] b = new byte[length];
        in.read(b, 0, length);

        // WARNING: String constructor is deprecated.
        bytes = new String(b, 0, length);
    }

    /** Used during output serialization by ClassFile only. */
    public void serialize(DataOutputStream out)
        throws IOException{
        out.writeByte(tag);
        out.writeChar(length);

        // WARNING: Handling of String coversion here might be problematic.
        out.writeBytes(bytes);
    }

    public String toString(){
        return bytes;
    }
}

```

A26. ConstantValue_attribute.java

Convenience class for representing ConstantValue_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Attribute that allows for initialization of static variables in
 * classes. This attribute will only reside in a field_info struct.
 */

public final class ConstantValue_attribute extends attribute_info{

    public int constantvalue_index;

    public ConstantValue_attribute(ClassFile cf, int ani, int al, int cvi){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        constantvalue_index = cvi;
    }

    public ConstantValue_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        constantvalue_index = in.readChar();
    }

    public void serialize(DataOutputStream out)
        throws IOException{
        attribute_length = 2;
        super.serialize(out);
        out.writeChar(constantvalue_index);
    }
}

```

A27. cp_info.java

Convenience class for representing cp_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Represents the common interface of all constant pool parts
 * that all specific constant pool items must inherit from.
 *
 */
public abstract class cp_info{

    /** The type tag that signifies what kind of constant pool
     * item it is */
    public int tag;

    /** Used for serialization of the object back into a bytestream. */
    abstract void serialize(DataOutputStream out) throws IOException;

    /** Default constructor. Simply does nothing. */
    public cp_info() {}

    /** Constructor simply takes in the ClassFile as a reference to
     * its parent
     */
    public cp_info(ClassFile cf) {}

    /** Used during input serialization by ClassFile only. */
    cp_info(ClassFile cf, DataInputStream in)
        throws IOException{
        tag = in.readUnsignedByte();
    }

}

```

A28. Deprecated_attribute.java

Convenience class for representing Deprecated_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** A fix attributed that can be located either in the ClassFile,
 * field_info or the method_info attribute. Mark deprecated to
 * indicate that the method, class or field has been superceded.
 */
public final class Deprecated_attribute extends attribute_info{

    public Deprecated_attribute(ClassFile cf, int ani, int al){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
    }

    /** Used during input serialization by ClassFile only. */
    Deprecated_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
    }

}

```

A29. Exceptions_attribute.java

Convenience class for representing Exceptions_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** This is the struct where the exceptions table are located.

```



```

* <br><br>
* This attribute can only appear once in a method_info struct.
*/
public final class Exceptions_attribute extends attribute_info{

    public int number_of_exceptions;
    public int[] exception_index_table;

    public Exceptions_attribute(ClassFile cf, int ani, int al, int noe,
                               int[] eit){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        number_of_exceptions = noe;
        exception_index_table = eit;
    }

    /** Used during input serialization by ClassFile only. */
    Exceptions_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        number_of_exceptions = in.readChar();
        exception_index_table = new int[number_of_exceptions];
        for (int i=0; i<number_of_exceptions; i++)
            exception_index_table[i] = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    public void serialize(DataOutputStream out)
        throws IOException{
        attribute_length = 2 + (number_of_exceptions*2);
        super.serialize(out);
        out.writeChar(number_of_exceptions);
        for (int i=0; i<number_of_exceptions; i++)
            out.writeChar(exception_index_table[i]);
    }
}

```

A30. field_info.java

Convenience class for representing field_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Represents the field_info structure as specified in the JVM
specification.
*/
public final class field_info{

    public int access_flags;
    public int name_index;
    public int descriptor_index;
    public int attributes_count;
    public attribute_info[] attributes;

    /** Convenience constructor. */
    public field_info(ClassFile cf, int flags, int ni, int di){
        access_flags = flags;
        name_index = ni;
        descriptor_index = di;
        attributes_count = 0;
        attributes = new attribute_info[0];
    }

    /** Constructor called only during the serialization process.
    * <br><br>

```

```

* This is intentionally left as package protected as we
* should not normally call this constructor directly.
* <br><br>
* Warning: the handling of len is not correct (after String s =...)
*/
field_info(ClassFile cf, DataInputStream in)
throws IOException{
    access_flags = in.readChar();
    name_index = in.readChar();
    descriptor_index = in.readChar();
    attributes_count = in.readChar();
    attributes = new attribute_info[attributes_count];
    for (int i=0; i<attributes_count; i++){
        in.mark(2);
        String s = cf.constant_pool[in.readChar()].toString();
        in.reset();
        if (s.equals("ConstantValue"))
            attributes[i] = new ConstantValue_attribute(cf, in);
        else if (s.equals("Synthetic"))
            attributes[i] = new Synthetic_attribute(cf, in);
        else if (s.equals("Deprecated"))
            attributes[i] = new Deprecated_attribute(cf, in);
        else
            attributes[i] = new attribute_info.Unknown(cf, in);
    }
}

/** To serialize the contents into the output format.
*/
public void serialize(DataOutputStream out)
throws IOException{
    out.writeChar(access_flags);
    out.writeChar(name_index);
    out.writeChar(descriptor_index);
    out.writeChar(attributes_count);
    for (int i=0; i<attributes_count; i++)
        attributes[i].serialize(out);
}
}

```

A31. InnerClasses_attribute.java

Convenience class for representing InnerClasses_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** A variable length structure that contains information about an
 * inner class of this class.
 */
public final class InnerClasses_attribute extends attribute_info{

    public int number_of_classes;
    public classes[] classes;

    public final static class classes{
        int inner_class_info_index;
        int outer_class_info_index;
        int inner_name_index;
        int inner_class_access_flags;
    }

    public InnerClasses_attribute(ClassFile cf, int ani, int al,
                                int noc, classes[] c){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
    }
}

```

```

        number_of_classes = noc;
        classes = c;
    }

    /** Used during input serialization by ClassFile only. */
    InnerClasses_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        number_of_classes = in.readChar();
        classes = new InnerClasses_attribute.classes[number_of_classes];
        for (int i=0; i<number_of_classes; i++){
            classes[i] = new classes();
            classes[i].inner_class_info_index = in.readChar();
            classes[i].outer_class_info_index = in.readChar();
            classes[i].inner_name_index = in.readChar();
            classes[i].inner_class_access_flags = in.readChar();
        }
    }

    /** Used during output serialization by ClassFile only. */
    public void serialize(DataOutputStream out)
        throws IOException{
        attribute_length = 2 + (number_of_classes * 8);
        super.serialize(out);
        out.writeChar(number_of_classes);
        for (int i=0; i<number_of_classes; i++){
            out.writeChar(classes[i].inner_class_info_index);
            out.writeChar(classes[i].outer_class_info_index);
            out.writeChar(classes[i].inner_name_index);
            out.writeChar(classes[i].inner_class_access_flags);
        }
    }
}

```

A32. LineNumberTable_attribute.java

Convenience class for representing LineNumberTable_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Determines which line of the binary code relates to the
 * corresponding source code.
 */
public final class LineNumberTable_attribute extends attribute_info{

    public int line_number_table_length;
    public line_number_table[] line_number_table;

    public final static class line_number_table{
        int start_pc;
        int line_number;
    }

    public LineNumberTable_attribute(ClassFile cf, int ani, int al, int
    lnt1,
                                line_number_table[] lnt){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        line_number_table_length = lnt1;
        line_number_table = lnt;
    }
}

```

```

    /** Used during input serialization by ClassFile only. */
    LineNumberTable attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        line_number_table_length = in.readChar();
        line_number_table = new
    LineNumberTable attribute.line_number_table(line_number_table_length);
        for (int i=0; i<line_number_table_length; i++){
            line_number_table[i] = new line_number_table();
            line_number_table[i].start_pc = in.readChar();
            line_number_table[i].line_number = in.readChar();
        }
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        attribute_length = 2 + (line_number_table_length * 4);
        super.serialize(out);
        out.writeChar(line_number_table_length);
        for (int i=0; i<line_number_table_length; i++){
            out.writeChar(line_number_table[i].start_pc);
            out.writeChar(line_number_table[i].line_number);
        }
    }
}

```

A33. LocalVariableTable_attribute.java

Convenience class for representing LocalVariableTable_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** Used by debugger to find out how the source file line number is linked
 * to the binary code. It has many to one correspondence and is found in
 * the Code_attribute.
 */
public final class LocalVariableTable_attribute extends attribute_info{

    public int local_variable_table_length;
    public local_variable_table[] local_variable_table;

    public final static class local_variable_table{
        int start_pc;
        int length;
        int name_index;
        int descriptor_index;
        int index;
    }

    public LocalVariableTable_attribute(ClassFile cf, int ani, int al,
                                       int lvt1, local_variable_table[]
lvt){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        local_variable_table_length = lvt1;
        local_variable_table = lvt;
    }

    /** Used during input serialization by ClassFile only. */
    LocalVariableTable_attribute(ClassFile cf, DataInputStream in)
        throws IOException{

```

```

        super(cf, in);
        local_variable_table_length = in.readChar();
        local_variable_table = new
LocalVariableTable_attribute.local_variable_table(local_variable_table_length);
        for (int i=0; i<local_variable_table_length; i++){
            local_variable_table[i] = new local_variable_table();
            local_variable_table[i].start_pc = in.readChar();
            local_variable_table[i].length = in.readChar();
            local_variable_table[i].name_index = in.readChar();
            local_variable_table[i].descriptor_index = in.readChar();
            local_variable_table[i].index = in.readChar();
        }
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{
        attribute_length = 2 + (local_variable_table_length * 10);
        super.serialize(out);
        out.writeChar(local_variable_table_length);
        for (int i=0; i<local_variable_table_length; i++){
            out.writeChar(local_variable_table[i].start_pc);
            out.writeChar(local_variable_table[i].length);
            out.writeChar(local_variable_table[i].name_index);
            out.writeChar(local_variable_table[i].descriptor_index);
            out.writeChar(local_variable_table[i].index);
        }
    }
}

```

A34. method_info.java

Convenience class for representing method_info structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** This follows the method_info in the JVM specification.
 */
public final class method_info {

    public int access_flags;
    public int name_index;
    public int descriptor_index;
    public int attributes_count;
    public attribute_info[] attributes;

    /** Constructor. Creates a method_info, initializes it with
     * the flags set, and the name and descriptor indexes given.
     * A new uninitialized code attribute is also created, and stored
     * in the <i>code</i> variable.*/
    public method_info(ClassFile cf, int flags, int ni, int di,
        int ac, attribute_info[] a) {
        access_flags = flags;
        name_index = ni;
        descriptor_index = di;
        attributes_count = ac;
        attributes = a;
    }

    /** This method creates a method_info from the current pointer in the
     * data stream. Only called by during the serialization of a complete
     * ClassFile from a bytestream, not normally invoked directly.
     */
    method_info(ClassFile cf, DataInputStream in)
        throws IOException{

```

```

        access_flags = in.readChar();
        name_index = in.readChar();
        descriptor_index = in.readChar();
        attributes_count = in.readChar();
        attributes = new attribute_info[attributes_count];
        for (int i=0; i<attributes_count; i++){
            in.mark(2);
            String s = cf.constant_pool[in.readChar()].toString();
            in.reset();
            if (s.equals("Code"))
                attributes[i] = new Code_attribute(cf, in);
            else if (s.equals("Exceptions"))
                attributes[i] = new Exceptions_attribute(cf, in);
            else if (s.equals("Synthetic"))
                attributes[i] = new Synthetic_attribute(cf, in);
            else if (s.equals("Deprecated"))
                attributes[i] = new Deprecated_attribute(cf, in);
            else
                attributes[i] = new attribute_info.Unknown(cf, in);
        }
    }

    /** Output serialization of the method_info to a byte array.
     * Not normally invoked directly.
     */
    public void serialize(DataOutputStream out)
        throws IOException{
        out.writeChar(access_flags);
        out.writeChar(name_index);
        out.writeChar(descriptor_index);
        out.writeChar(attributes_count);
        for (int i=0; i<attributes_count; i++)
            attributes[i].serialize(out);
    }
}

```

A35. SourceFile_attribute.java

Convenience class for representing SourceFile_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** A SourceFile attribute is an optional fixed length attribute in
 * the attributes table. Only located in the ClassFile struct only
 * once.
 */
public final class SourceFile_attribute extends attribute_info{

    public int sourcefile_index;

    public SourceFile_attribute(ClassFile cf, int ani, int al, int sfi){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
        sourcefile_index = sfi;
    }

    /** Used during input serialization by ClassFile only. */
    SourceFile_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
        sourcefile_index = in.readChar();
    }

    /** Used during output serialization by ClassFile only. */
    void serialize(DataOutputStream out)
        throws IOException{

```

```

        throws IOException{
        attribute_length = 2;
        super.serialize(out);
        out.writeChar(sourcefile_index);
    }
}

```

A36. Synthetic_attribute.java

Convenience class for representing Synthetic_attribute structures within ClassFiles.

```

import java.lang.*;
import java.io.*;

/** A synthetic attribute indicates that this class does not have
 *  a generated code source. It is likely to imply that the code
 *  is generated by machine means rather than coded directly. This
 *  attribute can appear in the classfile, method_info or field_info.
 *  It is fixed length.
 */
public final class Synthetic_attribute extends attribute_info{

    public Synthetic_attribute(ClassFile cf, int ani, int al){
        super(cf);
        attribute_name_index = ani;
        attribute_length = al;
    }

    /** Used during output serialization by ClassFile only. */
    Synthetic_attribute(ClassFile cf, DataInputStream in)
        throws IOException{
        super(cf, in);
    }
}

```

Annexure D1

Method void run()

```

0  getstatic #2 <Field java.lang.Object LOCK>
3  dup
4  astore_1
5  monitorenter
6  getstatic #3 <Field int counter>
9  iconst_1
10 iadd
11 putstatic #3 <Field int counter>
14 aload_1
15 monitorexit
16 return

```

Annexure D2

Method void run()

```

0  getstatic #2 <Field java.lang.Object LOCK>
3  dup
4  astore_1
5  dup
6  monitorenter
7  invokestatic #23 <Method void acquireLock(java.lang.Object)>
10 getstatic #3 <Field int counter>
13 iconst_1

```

```

14 iadd
15 putstatic #3 <Field int counter>
18 aload_1
19 dup
20 invokestatic #24 <Method void releaseLock(java.lang.Object)>
23 monitorexit
24 return

```

Annexure D3

```

import java.lang.*;

public class example{

    /** Shared static field. */
    public final static Object LOCK = new Object();

    /** Shared static field. */
    public static int counter = 0;

    /** Example method using synchronization. This method serves to
        illustrate the use of synchronization to implement thread-safe
        modification of a shared memory location by potentially multiple
        threads. */
    public void run(){

        // First acquire the lock, otherwise any memory writes we do will be
        // prone to race-conditions.
        synchronized (LOCK){

            // Now that we have acquired the lock, we can safely modify memory
            // in a thread-safe manner.
            counter++;

        }

    }

}

```

Annexure D4

```

import java.lang.*;
import java.util.*;
import java.net.*;
import java.io.*;

public class LockClient{

```



```

/** Protocol specific values. */
public final static int CLOSE = -1;
public final static int NACK = 0;
public final static int ACK = 1;
public final static int ACQUIRE_LOCK = 10;
public final static int RELEASE_LOCK = 20;

/** LockServer network values. */
public final static String serverAddress =
    System.getProperty("LockServer_network_address");
public final static int serverPort =
    Integer.parseInt(System.getProperty("LockServer_network_port"));

/** Table of global ID's for local objects. (hashCode-to-globalID
    mappings) */
public final static Hashtable hashCodeToGlobalID = new Hashtable();

/** Called when an application is to acquire a lock. */
public static void acquireLock(Object o){

    // First of all, we need to resolve the globalID for object 'o'.
    // To do this we use the hashCodeToGlobalID table.
    int globalID = ((Integer) hashCodeToGlobalID.get(o)).intValue();
    try{

        // Next, we want to connect to the LockServer, which will grant us
        // the global lock.
        Socket socket = new Socket(serverAddress, serverPort);
        DataOutputStream out =
            new DataOutputStream(socket.getOutputStream());
        DataInputStream in = new DataInputStream(socket.getInputStream());

        // Ok, now send the serialized request to the lock server.
        out.writeInt(ACQUIRE_LOCK);
        out.writeInt(globalID);
        out.flush();

        // Now wait for the reply.
        int status = in.readInt();           // This is a blocking call. So we
                                           // will wait until the remote side
                                           // sends something.

        if (status == NACK){
            throw new AssertionError(
                "Negative acknowledgement. Request failed.");
        }
    }
}

```

```
}else if (status != ACK){
    throw new AssertionError("Unknown acknowledgement: "
        + status + ". Request failed.");
}

// Close down the connection.
out.writeInt(CLOSE);
out.flush();
out.close();
in.close();

socket.close(); // Make sure to close the socket.

// This is a good acknowledgement, thus we can return now because
// global lock is now acquired.
return;

}catch (IOException e){
    throw new AssertionError("Exception: " + e.toString());
}

)

/** Called when an application is to release a lock. */
public static void releaseLock(Object o){

    // First of all, we need to resolve the globalID for object 'o'.
    // To do this we use the hashCodeToGlobalID table.
    int globalID = ((Integer) hashCodeToGlobalID.get(o)).intValue();

    try{

        // Next, we want to connect to the LockServer, which records us as
        // the owner of the global lock for object 'o'.
        Socket socket = new Socket(serverAddress, serverPort);
        DataOutputStream out =
            new DataOutputStream(socket.getOutputStream());
        DataInputStream in = new DataInputStream(socket.getInputStream());

        // Ok, now send the serialized request to the lock server.
        out.writeInt(RELEASE_LOCK);
        out.writeInt(globalID);
        out.flush();

        // Now wait for the reply.
```

```
int status = in.readInt();    // This is a blocking call. So we
                              // will wait until the remote side
                              // sends something.

if (status == NACK){
    throw new AssertionError(
        "Negative acknowledgement. Request failed.");
}else if (status != ACK){
    throw new AssertionError("Unknown acknowledgement: "
        + status + ". Request failed.");
}

// Close down the connection.
out.writeInt(CLOSE);
out.flush();
out.close();
in.close();

socket.close();              // Make sure to close the socket.

// This is a good acknowledgement, return because global lock is
// now released.
return;

}catch (IOException e){
    throw new AssertionError("Exception: " + e.toString());
}

}
```

Annexure D5

```
import java.lang.*;
import java.util.*;
import java.net.*;
import java.io.*;

public class LockServer implements Runnable{

    /** Protocol specific values */
    public final static int CLOSE = -1;
    public final static int NACK = 0;
```

```
public final static int ACK = 1;
public final static int ACQUIRE_LOCK = 10;
public final static int RELEASE_LOCK = 20;

/** LockServer network values. */
public final static int serverPort = 20001;

/** Table of lock records. */
public final static Hashtable locks = new Hashtable();

/** Linked list of waiting LockManager objects. */
public LockServer next = null;

/** Address of remote LockClient. */
public final String address;

/** Private input/output objects. */
private Socket socket = null;
private DataOutputStream outputStream;
private DataInputStream inputStream;

public static void main(String[] s)
throws Exception{

    System.out.println("LockServer_network_address="
        + InetAddress.getLocalHost().getHostAddress());
    System.out.println("LockServer_network_port=" + serverPort);

    // Create a serversocket to accept incoming lock operation
    // connections.
    ServerSocket serverSocket = new ServerSocket(serverPort);

    while (!Thread.interrupted()){

        // Block until an incoming lock operation connection.
        Socket socket = serverSocket.accept();

        // Create a new instance of LockServer to manage this lock
        // operation connection.
        new Thread(new LockServer(socket)).start();

    }

}
```

```
/** Constructor. Initialise this new LockServer instance with necessary
    resources for operation. */
public LockServer(Socket s){
    socket = s;
    try{
        outputStream = new DataOutputStream(s.getOutputStream());
        inputStream = new DataInputStream(s.getInputStream());
        address = s.getInetAddress().getHostAddress();
    }catch (IOException e){
        throw new AssertionError("Exception: " + e.toString());
    }
}

/** Main code body. Decode incoming lock operation requests and execute
    accordingly. */
public void run(){

    try{

        // All commands are implemented as 32bit integers.
        // Legal commands are listed in the "protocol specific values"
        // fields above.
        int command = inputStream.readInt();

        // Continue processing commands until a CLOSE operation.
        while (command != CLOSE){

            if (command == ACQUIRE_LOCK){                // This is an
                                                            // ACQUIRE_LOCK
                                                            // operation.

                // Read in the globalID of the object to be locked.
                int globalID = inputStream.readInt();

                // Synchronize on the locks table in order to ensure thread-
                // safety.
                synchronized (locks){

                    // Check for an existing owner of this lock.
                    LockServer lock = (LockServer) locks.get(
                        new Integer(globalID));

                    if (lock == null){                // No-one presently owns this lock,
                                                        // so acquire it.

```

```
        locks.put(new Integer(globalID), this);

        acquireLock();    // Signal to the client the
                          // successful acquisition of this
                          // lock.

    }else{                // Already owned. Append ourselves
                          // to end of queue.

        // Search for the end of the queue. (Implemented as
        // linked-list)
        while (lock.next != null){
            lock = lock.next;
        }

        lock.next = this; // Append this lock request at end.
    }

}

}else if (command == RELEASE_LOCK){    // This is a
                                      // RELEASE_LOCK
                                      // operation.

    // Read in the globalID of the object to be locked.
    int globalID = inputStream.readInt();

    // Synchronize on the locks table in order to ensure thread-
    // safety.
    synchronized (locks){

        // Check to make sure we are the owner of this lock.
        LockServer lock = (LockServer) locks.get(
            new Integer(globalID));

        if (lock == null){
            throw new AssertionError("Unlocked. Release failed.");
        }else if (lock.address != this.address){
            throw new AssertionError("Trying to release a lock "
                + "which this client doesn't own. Release "
                + "failed.");
        }

        lock = lock.next;
        lock.acquireLock();    // Signal to the client the
```

```
// successful acquisition of this
// lock.

// Shift the linked list of pending acquisitions forward
// by one.
locks.put(new Integer(globalID), lock);

// Clear stale reference.
next = null;
}

releaseLock(); // Signal to the client the successful
               // release of this lock.

}else{ // Unknown command.
    throw new AssertionError(
        "Unknown command. Operation failed.");
}

// Read in the next command.
command = inputStream.readInt();

}

} catch (Exception e){
    throw new AssertionError("Exception: " + e.toString());
}finally{
    try{
        // Closing down. Cleanup this connection.
        outputStream.flush();
        outputStream.close();
        inputStream.close();
        socket.close();
    }catch (Throwable t){
        t.printStackTrace();
    }
    // Garbage these references.
    outputStream = null;
    inputStream = null;
    socket = null;
}

}

/** Send a positive acknowledgement of an ACQUIRE_LOCK operation. */
```

```

        public void acquireLock() throws IOException{
            outputStream.writeInt(ACK);
            outputStream.flush();
        }

        /** Send a positive acknowledgement of a RELEASE_LOCK operation. */
        public void releaseLock() throws IOException{
            outputStream.writeInt(ACK);
            outputStream.flush();
        }
    }
}

```

ANNEXURE D6

LockLoader.java

This excerpt is the source-code of LockLoader, which modifies an application as it is being loaded.

```

import java.lang.*;
import java.io.*;
import java.net.*;

public class LockLoader extends URLClassLoader{

    public LockLoader(URL[] urls){
        super(urls);
    }

    protected Class findClass(String name)
    throws ClassNotFoundException{

        ClassFile cf = null;

        try{
            BufferedInputStream in =
                new BufferedInputStream(findResource(name.replace('.',
                    '/').concat(".class")).openStream());

            cf = new ClassFile(in);

        }catch (Exception e){throw new ClassNotFoundException(e.toString());}

        // Class-wide pointers to the enterindex and exitindex.
        int enterindex = -1;
        int exitindex = -1;

        for (int i=0; i<cf.methods_count; i++){
            for (int j=0; j<cf.methods[i].attributes_count; j++){
                if (!(cf.methods[i].attributes[j] instanceof Code_attribute))
                    continue;

                Code_attribute ca = (Code_attribute)
                    cf.methods[i].attributes[j];

                boolean changed = false;

                for (int z=0; z<ca.code.length; z++){
                    if ((ca.code[z][0] & 0xf) == 194){ // Opcode for a
                                                                // MONITORENTER

```



```

changed = true;

// instruction.

// Next, realign the code array, making room for the
// insertions.
byte[][] code2 = new byte[ca.code.length+2][];
System.arraycopy(ca.code, 0, code2, 0, z);
code2[z+1] = ca.code[z];
System.arraycopy(ca.code, z+1, code2, z+3,
    ca.code.length-(z+1));
ca.code = code2;

// Next, insert the DUP instruction.
ca.code[z] = new byte[1];
ca.code[z][0] = (byte) 89;

// Finally, insert the INVOKESTATIC instruction.
if (enterindex == -1){
    // This is the first time this class is encountering
    // acquirelock instruction, so have to add it to the
    // constant pool.
    cp_info[] cpi = new cp_info[cf.constant_pool.length+6];
    System.arraycopy(cf.constant_pool, 0, cpi, 0,
        cf.constant_pool.length);
    cf.constant_pool = cpi;
    cf.constant_pool_count += 6;

    CONSTANT_Utf8_info u1 =
        new CONSTANT_Utf8_info("LockClient");
    cf.constant_pool[cf.constant_pool.length-6] = u1;

    CONSTANT_Class_info c1 = new CONSTANT_Class_info(
        cf.constant_pool_count-6);
    cf.constant_pool[cf.constant_pool.length-5] = c1;

    u1 = new CONSTANT_Utf8_info("acquireLock");
    cf.constant_pool[cf.constant_pool.length-4] = u1;

    u1 = new CONSTANT_Utf8_info("(Ljava/lang/Object;)V");
    cf.constant_pool[cf.constant_pool.length-3] = u1;

    CONSTANT_NameAndType_info n1 =
        new CONSTANT_NameAndType_info(
            cf.constant_pool.length-4, cf.constant_pool.length-
3);
    cf.constant_pool[cf.constant_pool.length-2] = n1;

    CONSTANT_Methodref_info m1 = new
CONSTANT_Methodref_info(
2);
        cf.constant_pool.length-5, cf.constant_pool.length-
        cf.constant_pool.length-1] = m1;
        enterindex = cf.constant_pool.length-1;
    }
    ca.code[z+2] = new byte[3];
    ca.code[z+2][0] = (byte) 184;
    ca.code[z+2][1] = (byte) ((enterindex >> 8) & 0xff);
    ca.code[z+2][2] = (byte) (enterindex & 0xff);

    // And lastly, increase the CODE_LENGTH and
    // values.
    ca.code_length += 4;
    ca.attribute_length += 4;

    z += 1;

```

```

    } else if ((ca.code[z][0] & 0xff) == 195){ // Opcode for a
                                                // MONITOREXIT
                                                // instruction.

        changed = true;

        // Next, realign the code array, making room for the
        // insertions.
        byte[] code2 = new byte[ca.code.length+2][];
        System.arraycopy(ca.code, 0, code2, 0, z);
        code2[z+1] = ca.code[z];
        System.arraycopy(ca.code, z+1, code2, z+3,
            ca.code.length-(z+1));
        ca.code = code2;

        // Next, insert the DUP instruction.
        ca.code[z] = new byte[1];
        ca.code[z][0] = (byte) 89;

        // Finally, insert the INVOKESTATIC instruction.
        if (exitindex == -1){
            // This is the first time this class is encountering

            // acquirelock instruction, so have to add it to the
            // constant pool.
            cp_info[] cpi = new cp_info(cf.constant_pool.length+6);
            System.arraycopy(cf.constant_pool, 0, cpi, 0,
                cf.constant_pool.length);
            cf.constant_pool = cpi;
            cf.constant_pool_count += 6;

            CONSTANT_Utf8_info u1 =
                new CONSTANT_Utf8_info("LockClient");
            cf.constant_pool[cf.constant_pool.length-6] = u1;

            CONSTANT_Class_info c1 = new CONSTANT_Class_info(
                cf.constant_pool_count-6);
            cf.constant_pool[cf.constant_pool.length-5] = c1;

            u1 = new CONSTANT_Utf8_info("releaseLock");
            cf.constant_pool[cf.constant_pool.length-4] = u1;

            u1 = new CONSTANT_Utf8_info("(Ljava/lang/Object;)V");
            cf.constant_pool[cf.constant_pool.length-3] = u1;

            CONSTANT_NameAndType_info n1 =
                new CONSTANT_NameAndType_info(
                    cf.constant_pool.length-4, cf.constant_pool.length-
3);
            cf.constant_pool[cf.constant_pool.length-2] = n1;

            CONSTANT_Methodref_info m1 = new
CONSTANT_Methodref_info(
                cf.constant_pool.length-5, cf.constant_pool.length-
2);
                cf.constant_pool[cf.constant_pool.length-1] = m1;
                exitindex = cf.constant_pool.length-1;
            }
            ca.code[z+2] = new byte[3];
            ca.code[z+2][0] = (byte) 184;
            ca.code[z+2][1] = (byte) ((exitindex >> 8) & 0xff);
            ca.code[z+2][2] = (byte) (exitindex & 0xff);

            // And lastly, increase the CODE_LENGTH and
            ATTRIBUTE_LENGTH
            // values.
            ca.code_length += 4;

```

```
        ca.attribute_length += 4;

        z += 1;
    }
}

// If we changed this method, then increase the stack size by
one.
if (changed){
    ca.max_stack++;        // Just to make sure.
}
}

try(
    ByteArrayOutputStream out = new ByteArrayOutputStream();
    cf.serialize(out);

    byte[] b = out.toByteArray();

    return defineClass(name, b, 0, b.length);
}catch (Exception e){
    throw new ClassNotFoundException(name);
}
}
}
```

CLAIMS

1. A multiple computer system having at least one application program running simultaneously on a plurality of computers interconnected by a communications network, wherein a like plurality of substantially identical objects are created, each in the corresponding computer and each having a substantially identical name, and said system including a lock means applicable to all said computers wherein any computer wishing to utilize a named object therein acquires an authorizing lock from said lock means which permits said utilization and which prevents all the other computers from utilizing their corresponding named object until said authorizing lock is relinquished.
2. The system as claimed in claim 1 wherein said lock means includes an acquire lock routine and a release lock routine, and both said routines are included in modifications made to said application program running on all said computers.
3. The system as claimed in claim 2 wherein said lock means further includes a shared table listing said named objects in use by any said computer, a lock status for each said object, and a queue of any pending lock acquisitions.
4. The system as claimed in claim 3 wherein said lock means is located within an additional computer not running said application program and connected to said communications network.
5. The system as claimed in claim 2 wherein each said application program is modified before, during, or after loading by inserting a finalization routine to modify each instance at which said application program no longer needs to refer to an object.
6. The system as claimed in claim 3 wherein the application program is modified in accordance with a procedure selected from the group of procedures consisting of re-compilation at loading, pre-compilation prior to loading, compilation prior to loading, just-in-time compilation, and re-compilation after loading and before execution of the relevant portion of application program.
7. The system as claimed in claim 2 wherein said modified application program is transferred to all said computers in accordance with a procedure selected

from the group consisting of master/slave transfer, branched transfer and cascaded transfer.

8. A plurality of computers interconnected via a communications link and operating at least one application program simultaneously wherein each said computer in operating said at least one application program utilizes an object only in local memory physically located in each said computer, the contents of the local memory utilized by each said computer is fundamentally similar but not, at each instant, identical, and every one of said computers has an acquire lock routine and a release lock routine which permit utilization of the local object only by one computer and each of the remainder of said plurality of computers is locked out of utilization of their corresponding object.
9. The plurality of computers as claimed in claim 8 wherein the local memory capacity allocated to the or each said application program is substantially identical and the total memory capacity available to the or each said application program is said allocated memory capacity.
10. The plurality of computers as claimed in claim 8 wherein all said distribution update means communicate via said communications link at a data transfer rate which is substantially less than the local memory read rate.
11. The plurality of computers as claimed in claim 8 wherein at least some of said computers are manufactured by different manufacturers and/or have different operating systems.
12. A method of running at least one application program on a plurality of computers simultaneously, said computers being interconnected by means of a communications network, said method comprising the steps of:
 - (i) creating a like plurality of substantially identical objects each in the corresponding computer and each having a substantially identical name, and
 - (ii) requiring any of said computers wishing to utilize a named object therein to acquire an authorizing lock which permits said utilization and which prevents all the other computers from utilizing their corresponding named object until said authorizing lock is relinquished.

13. A method as claimed in claim 12 including the further step of:
 - (iii) providing each said computer with a distributed run time means to communicate between said computers via said communications network.
14. A method as claimed in claim 13 including the further step of:
 - (iv) providing a shared table accessible by each said distributed run time means and in which is stored the identity of any computer which currently has to access an object, together with the identity of the object.
15. A method as claimed in claim 14 including the further step of:
 - (v) associating a counter means with said shared table, said counter means storing a count of the number of said computers which seek access to said object.
16. A method as claimed in claim 15 including the further step of:
 - (vi) providing an additional computer on which said shared program does not run and which hosts said shared table and counter, said additional computer being connected to said communications network.
17. A method of ensuring consistent synchronization of an application program to be run simultaneously on a plurality of computers interconnected via a communications network, said method comprising the steps of:
 - (i) scrutinizing said application program at, or prior to, or after loading to detect each program step defining an synchronization routine, and
 - (ii) modifying said synchronization routine to ensure utilization of an object by only one computer and preventing all the remaining computers from simultaneously utilizing their corresponding objects.
18. The method claimed in claim 17 wherein step (ii) comprises the steps of:
 - (iii) loading and executing said synchronization routine on one of said computers,
 - (iv) modifying said synchronization routine by said one computer, and
 - (v) transferring said modified synchronization routine to each of the remaining computers.
19. The method as claimed in claim 18 wherein said modified synchronization routine is supplied by said one computer direct to each of said remaining computers.

20. The method as claimed in claim 18 wherein said modified synchronization routine is supplied in cascade fashion from said one computer sequentially to each of said remaining computers.
21. The method claimed in claim 17 wherein step (ii) comprises the steps of:
 - (vi) loading and modifying said synchronization routine on one of said computers,
 - (vii) said one computer sending said unmodified synchronization routine to each of the remaining computers, and
 - (viii) each of said remaining computers modifying said synchronization routine after receipt of same.
22. The method claimed in claim 21 wherein said unmodified synchronization routine is supplied by said one computer directly to each of said remaining computers.
23. The method claimed in claim 21 wherein said unmodified synchronization routine is supplied in cascade fashion from said one computer sequentially to each of said remaining computers.
24. The method as claimed in claim 17 including the further step of:
 - (ix) modifying said application program utilizing a procedure selected from the group of procedures consisting of re-compilation at loading, pre-compilation prior to loading, compilation prior to loading, just-in-time compilation, and re-compilation after loading and before execution of the relevant portion of application program.
25. The method as claimed in claim 17 including the further step of:
 - (x) transferring the modified application program to all said computers utilizing a procedure selected from the group consisting of master/slave transfer, branched transfer and cascaded transfer.
26. In a multiple thread processing computer operation in which individual threads of a single application program are simultaneously being processed each on a corresponding one of a plurality of computers interconnected via a communications link, and in which objects in local memory physically associated with the computer processing each thread have corresponding objects in the local memory of each other said computer, the improvement

comprising permitting only one of said computers to utilize an object and preventing all the remaining computers from simultaneously utilizing their corresponding object.

27. The improvement as claimed in claim 26 wherein an object residing in the memory associated with one said thread and to be utilized has its identity communicated by the computer of said one thread to a shared table accessible by all other said computers.
28. The improvement as claimed in claim 26 wherein an object residing in the memory associated with one said thread and to be utilized has its identity transmitted to the computer associated with another said thread and is transmitted thereby to a shared table accessible by all said other computers.
29. A computer program product comprising a set of program instructions stored in a storage medium and operable to permit a plurality of computers to carry out the method as claimed in claim 12 or 17.
30. A plurality of computers interconnected via a communication network and operable to ensure consistent initialization of an application program running simultaneously of said computers, said computers being programmed to carry out the method as claimed in claim 12 or 17 or being loaded with the computer program product as claimed in claim 29.

1/15

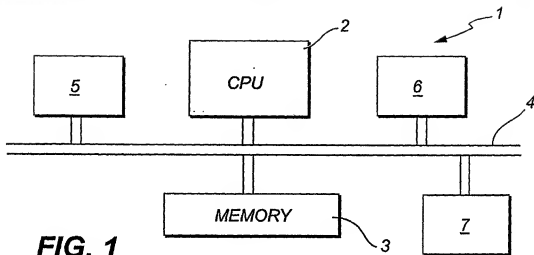


FIG. 1
PRIOR ART

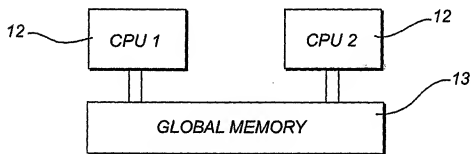


FIG. 2
PRIOR ART

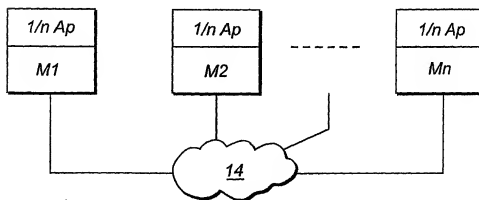


FIG. 3
PRIOR ART

2/15

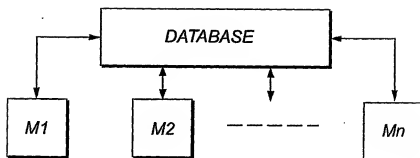


FIG. 4
PRIOR ART

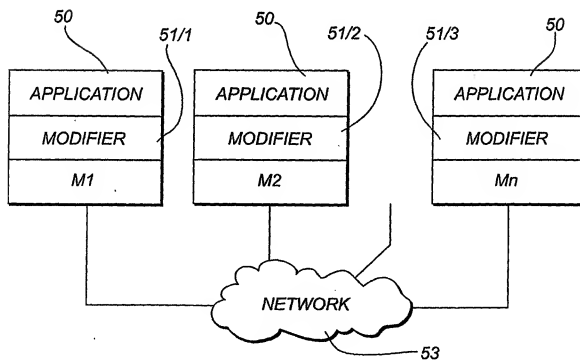
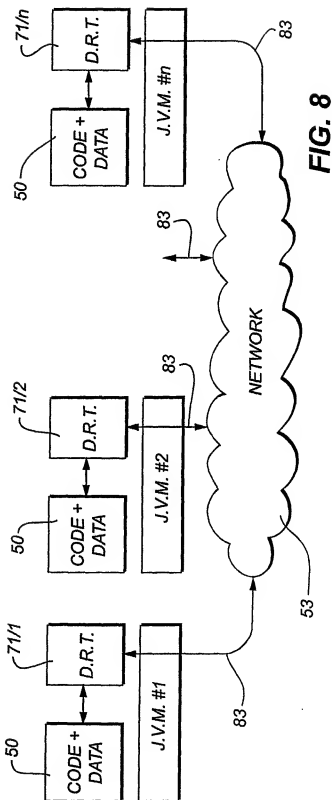
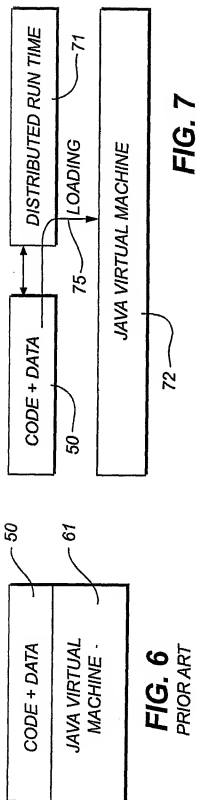
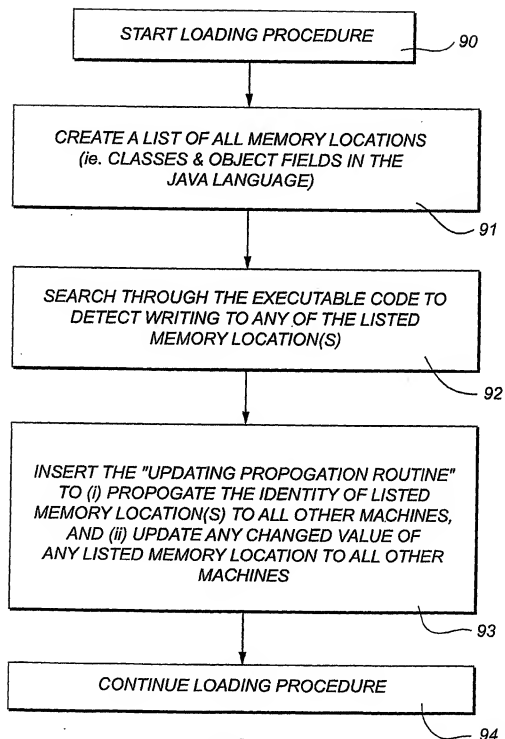


FIG. 5

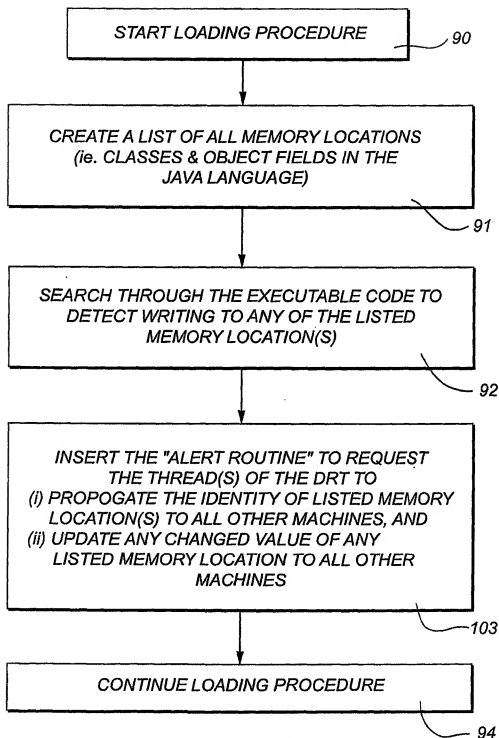
3/15



4/15

**FIG. 9**

5/15

**FIG. 10**

6/15

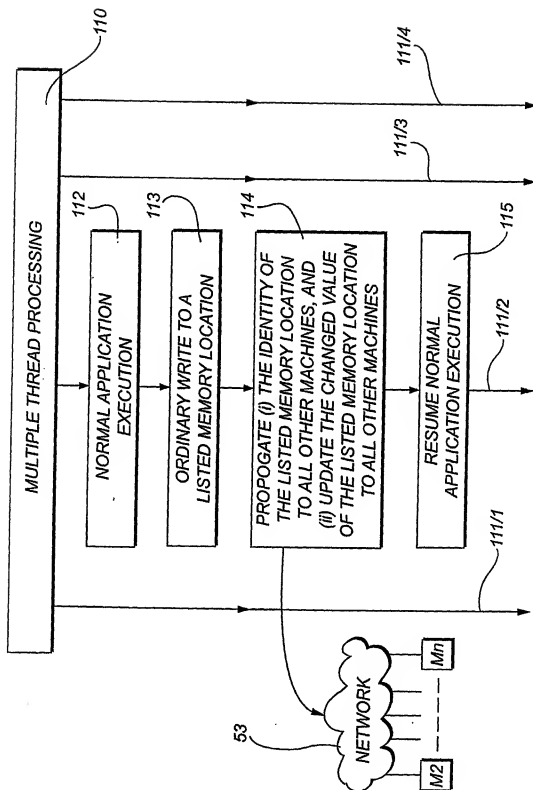


FIG. 11

7/15

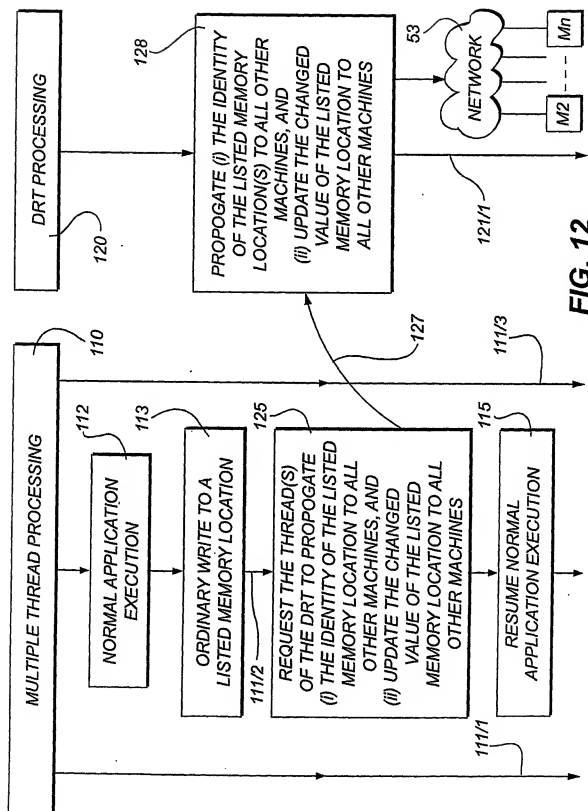
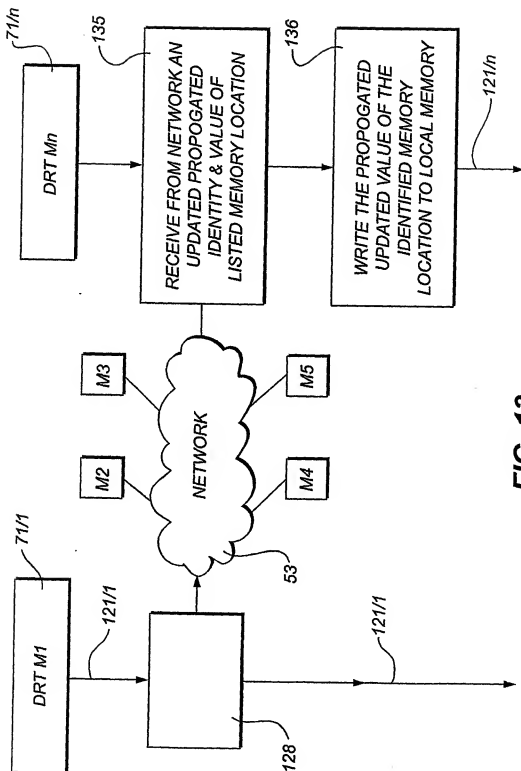


FIG. 12

**FIG. 13**

9/15

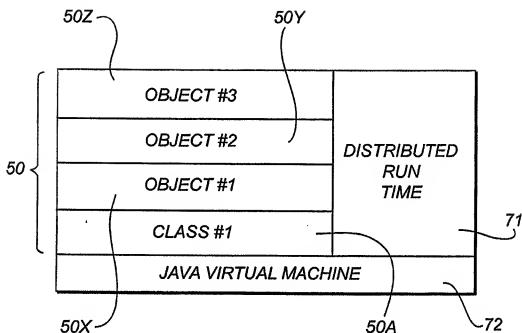
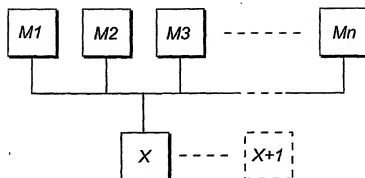
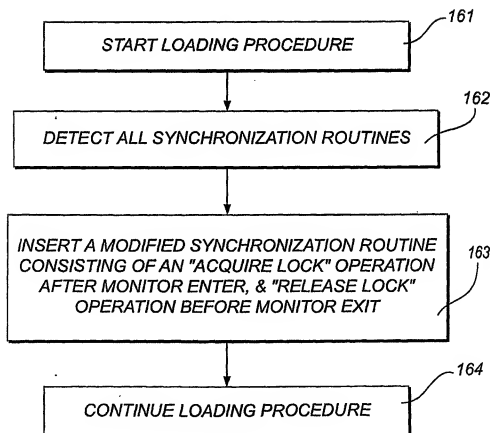
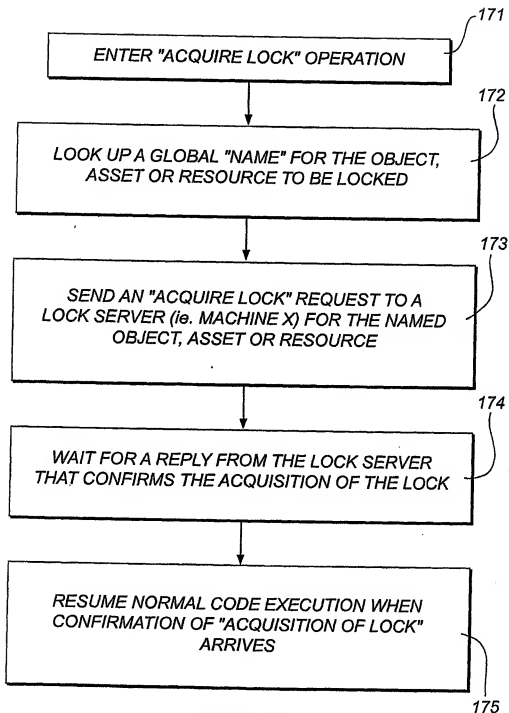


FIG. 14
PRIOR ART

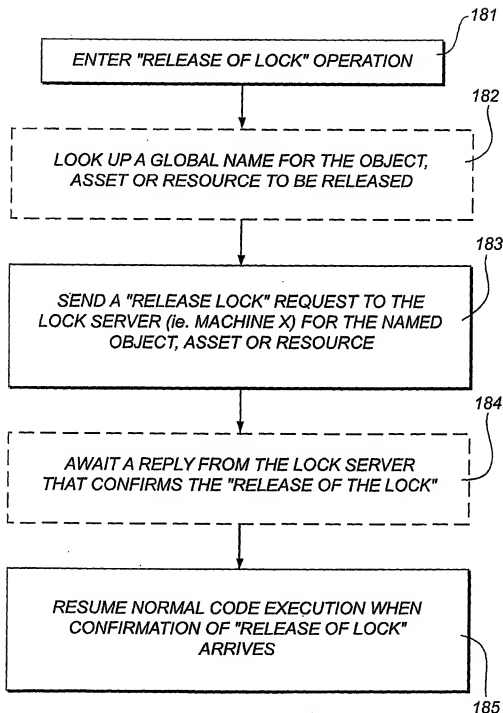
10/15

**FIG. 15****FIG. 16**

11/15

**FIG. 17**

12/15

**FIG. 18**

13/15

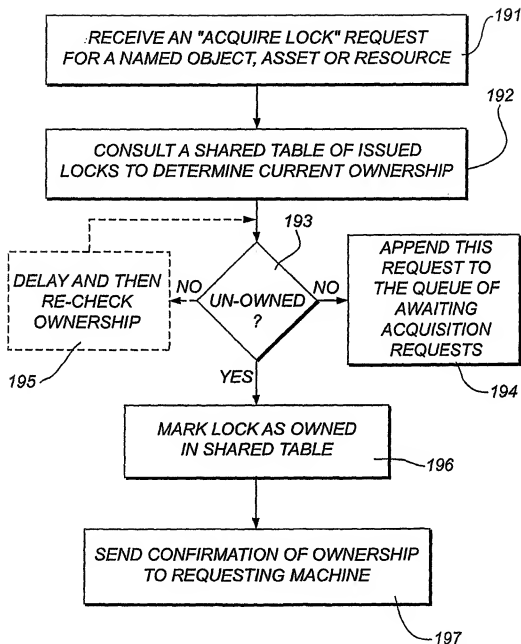
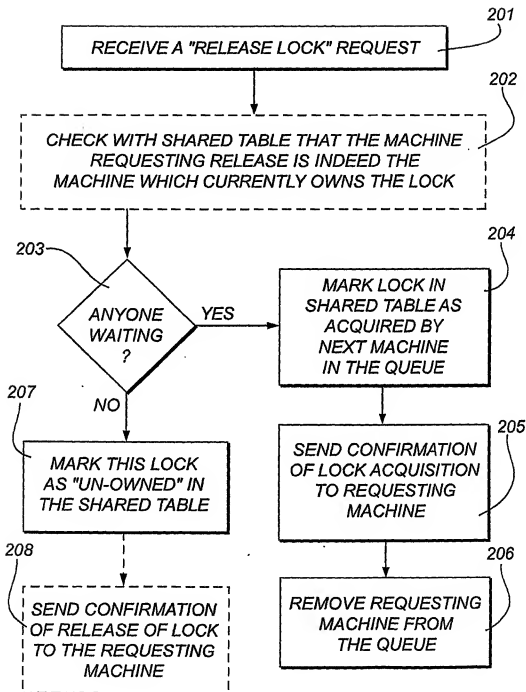
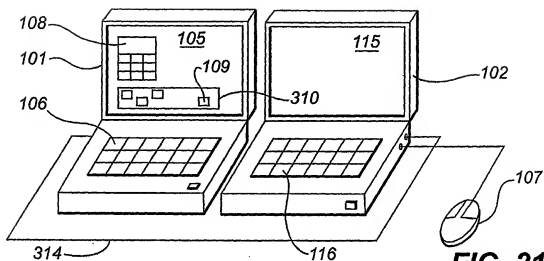
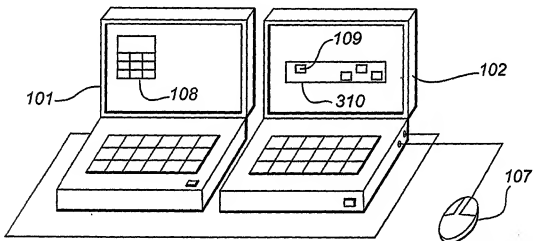
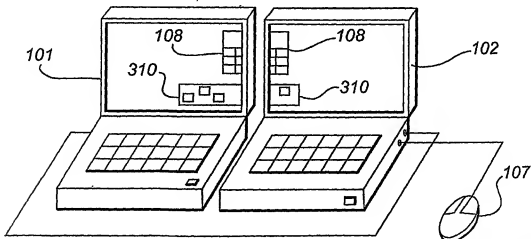


FIG. 19

14/15

**FIG. 20**

15/15

**FIG. 21****FIG. 22****FIG. 23**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2005/000579

A. CLASSIFICATION OF SUBJECT MATTER		
Int. Cl. ⁷ : G06F 15/16		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols)		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) DWPI, USPTO, PCT, IEEE, internet (object, replicate, redundant, mirror, cluster, distributed, lock, thread, compile, high availability, runtime, virtual machine, etc.)		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2003/083614 A1 (ETERNAL SYSTEMS, INC.), 9 October 2003 the whole document	1-30
X	US 6,625,751 B1 (STAROVIC et al), 23 September 2003 the whole document	1-30
X	WO 2002/044835 A2 (GINGERICH), 6 June 2002 the whole document	1-30
X	US 2004/0073828 A1 (BRONSTEIN), 15 April 2004 the whole document	1-30
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input checked="" type="checkbox"/> See patent family annex		
* Special categories of cited documents:		
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention	
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art	
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family	
"P" document published prior to the international filing date but later than the priority date claimed		
Date of the actual completion of the international search 4 July 2005	Date of mailing of the international search report 11 JUL 2005	
Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA E-mail address: pct@ipaustralia.gov.au Facsimile No. (02) 6285 3929	Authorized officer MATTHEW HOLLINGWORTH Telephone No : (02) 6283 2024	

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2005/000579

C (Continuation).

DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,488,723 A (BARADEL et al), 30 January 1996 the whole document	1-30
X	T. C. Bressoud, <i>TFT: A Software System for Application-Transparent Fault Tolerance</i> Proc. 28 th Annual International Symposium on Fault-Tolerant Computing, pp. 128-37, 1998	1-30

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2005/000579

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report		Patent Family Member			
WO 03083614	AU	2003223352	AU	2003230748	EP 1495414
	EP	1495571	US	2004078617	US 2004078618
	WO	03084116			
US 6625751	GB	2353113			
WO 0244835	AU	32411/02			
US 2004073828	NONE				
US 5488723	EP	0572307	FR	2691559	WO 9324884
Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.					
END OF ANNEX					